arXiv:1604.04013v2 [cs.PF] 18 Jun 2016

Ergodic Theory for Controlled Markov Chains with Stationary Inputs

Yue Chen, Ana Bušić, and Sean Meyn[*†]

June 21, 2016

Abstract

Consider a stochastic process $\boldsymbol{X}$ on a finite state space $\mathsf{X} = \{1, \dots, d\}$. It is conditionally Markov, given a real-valued 'input process' $\boldsymbol{\zeta}$. This is assumed to be small, which is modeled through the scaling,

$$\zeta_t = \varepsilon \zeta^1_t, \qquad 0 \le \varepsilon \le 1\,,$$

where $\boldsymbol{\zeta}^1$ is a bounded stationary process. The following conclusions are obtained, subject to smoothness assumptions on the controlled transition matrix and a mixing condition on $\boldsymbol{\zeta}$:

(i) A stationary version of the process is constructed, that is coupled with a stationary version of the Markov chain $\boldsymbol{X}^\bullet$ obtained with $\boldsymbol{\zeta} \equiv 0$. The triple $(\boldsymbol{X}, \boldsymbol{X}^\bullet, \boldsymbol{\zeta})$ is a jointly stationary process satisfying

$$\mathsf{P}\{X(t) \neq X^\bullet(t)\} = O(\varepsilon)$$

Moreover, a second-order Taylor-series approximation is obtained:

$$\mathsf{P}\{X(t) = i\} = \mathsf{P}\{X^\bullet(t) = i\} + \varepsilon^2 \varrho(i) + o(\varepsilon^2), \quad 1 \le i \le d,$$

with an explicit formula for the vector $\varrho \in \mathbb{R}^d$.

(ii) For any $m \ge 1$ and any function $f\colon \{1, \dots, d\} \times \mathbb{R} \to \mathbb{R}^m$, the stationary stochastic process $Y(t) = f(X(t), \zeta(t))$ has a power spectral density S_f that admits a second order Taylor series expansion: A function $\mathrm{S}^{(2)}_f\colon [-\pi, \pi] \to \mathbb{C}^{m \times m}$ is constructed such that

$$\mathrm{S}_f(\theta) = \mathrm{S}^\bullet_f(\theta) + \varepsilon^2 \mathrm{S}^{(2)}_f(\theta) + o(\varepsilon^2), \quad \theta \in [-\pi, \pi].$$

An explicit formula for the function $\mathrm{S}^{(2)}_f$ is obtained, based in part on the bounds in (i).

The results are illustrated using a version of the timing channel of Anantharam and Verdu.

[*]This research is supported by the NSF grants CPS-0931416 and CPS-1259040, the French National Research Agency grant ANR-12-MONU-0019, and US-Israel BSF Grant 2011506.

[†]Y.C. and S.M. are with the Department of Electrical and Computer Engg. at the University of Florida, Gainesville. A.B. is with Inria and the Computer Science Dept. of École Normale Supérieure, Paris, France.

1 Introduction

This paper concerns second-order ergodic theory for a controlled Markov chain. Consider for the sake of illustration a stochastic process $\boldsymbol{X}$ on a finite state space $\mathsf{X} = \{1, \ldots, d\}$, which evolves together with a real-valued stationary sequence $\boldsymbol{\zeta}$ and an i.i.d. sequence $\boldsymbol{N}$ according to the recursion,

$$X_{t+1} = \varphi(X_t, \zeta_t, N_{t+1}), \qquad t \in \mathbb{Z} \tag{1}$$

where $\varphi \colon \mathsf{X} \times \mathbb{R}^2 \to \mathsf{X}$ is Borel measurable. In the special case $\boldsymbol{\zeta} \equiv 0$ we denote the solution $\boldsymbol{X}^\bullet$, which is a time-homogeneous Markov chain.

Among the questions we might ask are,

(i) If we are given only the function φ, the stationary process $\boldsymbol{\zeta}$, and the i.i.d. sequence $\boldsymbol{N}$, does there exist a stationary solution to (1)?

(ii) Can we compute statistics of $\boldsymbol{X}$, such as the marginal distribution?

(iii) If the answer to (ii) is no, can we *approximate* the statistics of $\boldsymbol{X}$? How do these statistics compare with the stationary Markov chain $\boldsymbol{X}^\bullet$?

Under the assumptions imposed in this paper, we construct the joint stationary process $(\boldsymbol{X}, \boldsymbol{X}^\bullet, \boldsymbol{\zeta})$ on the same probability space.

An answer to (ii) requires special assumptions on the model that are far beyond the scope of this paper. Instead we consider a setting in which $\boldsymbol{\zeta}$ is small, and obtain approximations to address question (iii). It is simplest to consider a family of processes, parameterized by a small constant $\varepsilon > 0$,

$$\zeta_t = \varepsilon \zeta_t^1,$$

where $\boldsymbol{\zeta}^1 = \{\zeta_t^1\}$ is a bounded sequence. The construction of $(\boldsymbol{X}, \boldsymbol{X}^\bullet, \boldsymbol{\zeta})$ is obtained so that for each t,

$$\mathsf{P}\{X_t \neq X_t^\bullet\} = O(\varepsilon) \tag{2}$$

This then provides tools to address (iii).

In order to apply techniques from second-order statistics, the process is lifted to the unit simplex in $\mathbb{R}^d$. We denote

$$\Gamma_t = e^j, \qquad \text{when } X_t = j, \tag{3}$$

where e^j denotes the jth standard basis element in $\mathbb{R}^d$. This is a standard construction; it is useful since the evolution of $\{\Gamma_t\}$ can be expressed as a linear state space model driven by an uncorrelated "noise process" (see (8) below). This linear representation is used in [1] to construct a Kalman filter for a time-homogeneous Markov chain (without the input $\boldsymbol{\zeta}$), and these results are extended to a class of controlled Markov chains in [2].

The initial motivation for [2], as well as the research described here, is application to distributed control for the purposes of "demand dispatch" using distributed resources in a power grid. The results of the present paper are applied in [3] to obtain performance approximations in the same power grid model. Similar bounds were previously obtained in [4], but this is the first paper to obtain an exact second-order Taylor series approximation for second-order statistics.

The main contribution of this paper is to obtain tight approximations for the joint auto-correlation function for $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$, and hence also its power spectral density. To obtain these results requires the coupling bound (2), a second order Taylor series expansion for $\pi^\varepsilon = \mathsf{E}[\Gamma_t]$ in steady-state, and surprisingly complex calculations for a linearized model.

We are not aware of other related results in the literature. However, in some of our approximations we borrow one technique from [5] – the use of the fundamental matrix appears in the approximation of π^{ε}; see (18) for a definition, and further explanation following this equation.

The main results are summarized in Section 2, with all of the technical proofs contained in an appendix. Section 3 contains numerical results for an application to information theory – a variant of the timing channel introduced in [6]. Conclusions and directions for future research are contained in Section 4.

2 Model and main results

Consider an irreducible and aperiodic Markov chain $\boldsymbol{X}^{\bullet}$ evolving on a finite state space $\mathsf{X} = \{1, \ldots, d\}$, and transition matrix P_0. This admits a stationary realization on the two sided time-interval $\mathbb{Z}$, whose marginal distribution π_0 is the invariant probability mass function (pmf) for P_0, satisfying $\pi_0 P_0 = \pi_0$. The goal of this paper is to investigate how the statistics change when the dynamics are subject to an exogenous disturbance.

2.1 Controlled Markov model

The stochastic process $\boldsymbol{X}$ considered in this paper also evolves on the finite state space X. The "disturbance" in the controlled model is a one-dimensional stationary process denoted $\boldsymbol{\zeta} = \{\zeta_t : -\infty < t < \infty\}$. A controlled transition matrix $\{P_\zeta : \zeta \in \mathbb{R}\}$ describes the dynamics of the process:

$$\mathsf{P}\{X_{t+1} = k \mid \zeta_s,\ X_s : s \le t\} = P_\zeta(j,k) \qquad \text{when } \zeta_t = \zeta,\ X_t = j. \tag{4}$$

It is assumed that P_ζ is a smooth function of ζ, and that P_0 is the transition matrix for $\boldsymbol{X}^{\bullet}$.

Since $\boldsymbol{X}$ is no longer Markovian, we cannot apply standard Markov chain theory to investigate properties of a stationary version of $\boldsymbol{X}$. Instead we apply linear systems theory, and for this we require a linear systems representation for the controlled stochastic process.

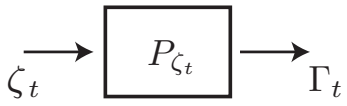

Figure 1: Controlled Markov model.

This is obtained by embedding the process in $\mathbb{R}^d$ through the indicator process Γ_t defined in (3). Linear dynamics are obtained by considering a specific realization of the model. We assume that there is a $d \times d$ matrix-valued function $\mathcal{G}$ and an i.i.d. sequence $\boldsymbol{N}$ for which,

$$\Gamma_{t+1} = \Gamma_t G_{t+1}, \quad G_{t+1} = \mathcal{G}(\zeta_t, N_{t+1}) \tag{5}$$

It is assumed moreover that $\boldsymbol{N}$ is independent of $\boldsymbol{\zeta}$, and that the entries of $\mathcal{G}$ are zero or one, with

$$\sum_{k=1}^{d} \mathcal{G}_{j,k}(z,n) = 1, \qquad \text{for all } j, z, n.$$

It follows from (4) that for each t,

$$\mathsf{E}[\mathcal{G}(\zeta_t, N_{t+1}) \mid \zeta_{-\infty}^{\infty}] = P_{\zeta_t}\,. \tag{6}$$

The random linear system (5) illustrated in Fig. 1 is the focus of study in this paper. The sequence $\boldsymbol{\Gamma} = \{\Gamma_t\}$ is viewed as a state process, that is driven by the disturbance (or "input") $\boldsymbol{\zeta}$. The state process evolves on the extreme points of the unit simplex in $\mathbb{R}^d$. We let $\boldsymbol{\Gamma}^{\bullet} = \{\Gamma_t^{\bullet}\}$ denote the stationary Markov chain obtained with $\boldsymbol{\zeta} \equiv 0$.

Our main assumptions are summarized in the following:

A1: The transition matrix P_0 is irreducible and aperiodic. The matrix valued function P_ζ is twice continuously differentiable (C^2) in a neighborhood of $\zeta = 0$, and the second derivative is Lipschitz continuous.

A2: $\zeta_t = \varepsilon \zeta_t^1$ where $\boldsymbol{\zeta}^1 = \{\zeta_t^1 : t \in \mathbb{Z}\}$ is a real-valued stationary stochastic process with zero mean. The following additional assumptions are imposed:

(i) It is bounded, $|\zeta_t^1| \le 1$ for all t with probability one. Hence $\sigma^2_{\zeta^1} = \mathsf{E}[(\zeta_t^1)^2] \le 1$.

(ii) Its auto-covariance is absolutely summable:

$$\sum_{t=0}^{\infty} |R_{\zeta^1}(t)| < \infty$$

The power spectral density S_{ζ^1} exists and is continuous under Assumption **A2** (ii). It also admits a spectral factor, denoted H_{ζ^1}:

$$\mathrm{S}_{\zeta^1}(\theta) = H_{\zeta^1}(e^{j\theta}) H_{\zeta^1}(e^{-j\theta}), \quad -\pi \le \theta \le \pi. \tag{7}$$

See [7] for background.

Assumption **A1** is used to obtain the approximation of (5) by an LTI (linear time invariant) system. The following intermediate step is used in [1] to obtain a Kalman filter for an uncontrolled Markov chain. The proof of Proposition 2.1 is given in Appendix B.

Proposition 2.1. *The random linear system* (5) *can be represented as*

$$\Gamma_{t+1} = \Gamma_t P_{\zeta_t} + \Delta_{t+1}, \tag{8}$$

where $\Delta_{t+1} = \Gamma_t(G_{t+1} - P_{\zeta_t})$. *This is a martingale difference sequence, with covariance matrix*

$$\Sigma^\Delta = Cov\left(\Gamma_t[G_{t+1} - P_{\zeta_t}]\right) = \mathsf{E}\big[\,\mathrm{diag}(\Gamma_{t+1}) - P_{\zeta_t}^T \Lambda_t^\Gamma P_{\zeta_t}\big] \tag{9}$$

where Λ_t^Γ *is the diagonal matrix with diagonal entries* $\{\Gamma_t(x^i) : 1 \le i \le d\}$. *Moreover,*

$$R_{\Delta,\zeta}(t) = 0, \quad \textit{for all } t. \tag{10}$$

That is, $\boldsymbol{\Delta}$ *and* $\boldsymbol{\zeta}$ *are uncorrelated.* □

We next apply the second-order Taylor series approximation:

$$P_{\zeta_t} = P_0 + \mathcal{E}\zeta_t + \frac{1}{2}\mathcal{W}\zeta_t^2 + O(\varepsilon^3)\,,$$

where $\mathcal{E}$ and $\mathcal{W}$ denote the first and second derivatives of P_ζ, evaluated at $\zeta = 0$:

$$\frac{d}{d\zeta} P_\zeta \bigg|_{\zeta=0} = \mathcal{E}, \qquad \frac{d^2}{d\zeta^2} P_\zeta \bigg|_{\zeta=0} = \mathcal{W}.$$

The $O(\varepsilon^3)$ bound holds under the Lipschitz condition for the second derivative of P_ζ.

The recursion (8) can be approximated as

$$\Gamma_{t+1} = \Gamma_t(P_0 + \mathcal{E}\zeta_t + \frac{1}{2}\mathcal{W}\zeta_t^2) + \Delta_{t+1} + O(\varepsilon^3).$$

This is the LTI approximation:

Proposition 2.2. *The system (5) can be represented by,*

$$\Gamma_{t+1} = \Gamma_t P_0 + D_{t+1} + O(\varepsilon^3)\,, \tag{11}$$

where, $D_{t+1} = B_t^T \zeta_t + V_t^T \zeta_t^2 + \Delta_{t+1}$, *with*

$$B_t^T = \Gamma_t \mathcal{E}, \qquad V_t^T = \tfrac{1}{2}\Gamma_t \mathcal{W}\,. \tag{12}$$

□

Applying the LTI approximation (11), an approximation for the auto-correlation of $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$ is obtained from an approximation for the pair process $(\boldsymbol{D}, \boldsymbol{\zeta})$. Since $\boldsymbol{D}$ is taken as a row vector, we use the following notation for the auto-correlation of $(\boldsymbol{D}, \boldsymbol{\zeta})$:

$$R(t) = \begin{bmatrix} R_D(t) & R_{D,\zeta}(t) \\ R_{D,\zeta}(-t)^T & R_\zeta(t) \end{bmatrix} \tag{13}$$

where $R_D(t) = \mathsf{E}[D(t)^T D(0)]$, $R_{D,\zeta}(t) = \mathsf{E}[D(t)^T \zeta(0)]$, and the expectations are taken in steady-state.

The *existence* of a steady-state solution is established in Prop. 2.3 that follows.

2.2 Correlation formulae and approximations

Under Assumptions we obtain a coupling result, which plays an important role in the approximations that follow. We write $\Gamma_t = \Gamma_t^\bullet + \widetilde{O}(\varepsilon)$ if

$$\mathsf{E}[\|\Gamma_t - \Gamma_t^\bullet\|] = O(\varepsilon)\,,$$

which implies that (2) also holds. We adopt similar notation for other random variables. The following result is proven in Appendix A:

Proposition 2.3. *Under Assumptions* ***A1*** *and* ***A2****, there exists* $\varepsilon_0 > 0$ *such that the following holds for each* $\varepsilon \in (0, \varepsilon_0]$*: the two process* $\boldsymbol{\Gamma}$ *and* $\boldsymbol{\Gamma}^\bullet$ *can be constructed so that* $(\boldsymbol{\Gamma}, \boldsymbol{\Gamma}^\bullet, \boldsymbol{\zeta})$ *is jointly stationary on the two-sided time interval* $\mathbb{Z}$, $\Gamma_t^\bullet$ *is independent of* $\boldsymbol{\zeta}$*, and moreover*

$$\Gamma_t = \Gamma_t^\bullet + \widetilde{O}(\varepsilon) \tag{14}$$

$$\mathsf{E}[\Gamma_t \zeta_t] = O(\varepsilon^2) \tag{15}$$

Consequently, for the stationary process,

$$B_t = B_t^\bullet + \widetilde{O}(\varepsilon), \quad V_t = V_t^\bullet + \widetilde{O}(\varepsilon), \quad \Delta_t = \Delta_t^\bullet + \widetilde{O}(\varepsilon) \tag{16}$$

□

The following strengthening of Assumption **A2** is useful in computations:

A3: The transfer function H_{ζ^1} in (7) is rational, with distinct poles $\{\rho_1, \dots, \rho_{n_z}\}$ satisfying $|\rho_i| < 1$ for each i.

Under **A2** and **A3** the auto-covariance function for $\boldsymbol{\zeta}$ can be expressed as a sum of geometrically decaying terms,

$$R_\zeta(t) = \varepsilon^2 \sum_{k=1}^{n_z} a_k \rho_k^{|t|}\,, \tag{17}$$

where the $\{a_k\}$ can be determined from H_{ζ^1}. Approximations for the auto-correlation functions $R_{D,\zeta}(t)$ and $R_D(t)$ in (13) are given in Theorem 2.4.

As in the perturbation theory of [5], one component in these approximations is based on the *fundamental matrix*,

$$U_1 = [I - P_0 + 1 \otimes \pi_0]^{-1} \tag{18}$$

where $1 \otimes \pi_0$ denotes the matrix whose rows are identical, and equal to π_0. Because the chain is irreducible and aperiodic, this can be expressed as a power series expansion,

$$U_1 = I + \sum_{k=1}^{\infty} [P_0 - 1 \otimes \pi_0]^k$$

The summand can also be expressed $[P_0 - 1 \otimes \pi_0]^k = [P_0^k - 1 \otimes \pi_0]$, $k \ge 1$. Hence convergence of the sum follows from the mean ergodic theorem,

$$\lim_{k\to\infty} P_0^k = 1 \otimes \pi_0 \tag{19}$$

where the rate of convergence is geometric.

Theorem 2.4. *Suppose that Assumptions **A1** and **A2** hold, and consider the stationary process $(\boldsymbol{\Gamma}, \boldsymbol{\Gamma}^\bullet, \boldsymbol{\zeta})$ constructed in Proposition 2.3, with $\varepsilon \in (0, \varepsilon_0]$. Then, for each t,*

$$R_{D,\zeta}(t) = BR_\zeta(t-1) + O(\varepsilon^3) \tag{20}$$

$$\begin{aligned} R_D(t) = {} & R_{B\zeta}(t) && \text{(21a)}\\ & + R_\Delta(t) && \text{(21b)}\\ & + R_{B\zeta,\Delta}(t-1) + R^T_{B\zeta,\Delta}(-t-1) && \text{(21c)}\\ & + R_{V\zeta^2,\Delta}(t-1) + R^T_{V\zeta^2,\Delta}(-t-1) && \text{(21d)}\\ & + O(\varepsilon^3) \end{aligned}$$

in which $B^T = \pi_0 \mathcal{E}$ in (20), and each component shown on the right hand side of (21a)–(21d) is given below:

(a) *The auto-correlation $R_{B\zeta}(t) = \mathsf{E}[B_t \zeta_t B_0^T \zeta_0]$ in (21a) admits the approximation,*

$$R_{B\zeta}(t) = (P_0^t \mathcal{E})^T \Pi_0 \mathcal{E} R_\zeta(t) + O(\varepsilon^3), \quad t \ge 0 \tag{22}$$

where $\Pi_0 = \operatorname{diag}(\pi_0)$.

(b) *The covariance for the martingale-difference sequence $\boldsymbol{\Delta}$ is given by $R_\Delta(t) = 0$ for $t \ne 0$, and*

$$\begin{aligned} R_\Delta(0) = \Sigma^\Delta = \Pi_\varepsilon & - P_0^T \Pi_\varepsilon P_0 \\ & - \left[P_0^T \operatorname{diag}(R_{\Gamma,\zeta}(0)) \mathcal{E} + \mathcal{E}^T \operatorname{diag}(R_{\Gamma,\zeta}(0)) P_0\right] \\ & - \tfrac{1}{2} R_\zeta(0) \left[P_0^T \Pi_0 \mathcal{W} + 2\mathcal{E}^T \Pi_0 \mathcal{E} + \mathcal{W}^T \Pi_0 P_0\right] \end{aligned} \tag{23}$$

where $\Pi_\varepsilon = \operatorname{diag}(\pi_\varepsilon)$, *with* $\pi_\varepsilon = \mathsf{E}[\Gamma_t]$, *and*

$$R_{\Gamma,\zeta}(t) = \mathsf{E}[\Gamma(t)^{\scriptscriptstyle T}\zeta(0)], \quad t \in \mathbb{Z}\,. \tag{24}$$

When $\varepsilon = 0$, *eq.* (23) *becomes*

$$\Sigma^{\Delta^\bullet} = \Pi_0 - P_0^{\scriptscriptstyle T}\Pi_0 P_0 \tag{25}$$

(c) *The cross-covariance* $R_{B\zeta,\Delta}(t) = \mathsf{E}[B_t\zeta_t\Delta_0]$ *admits the approximation,*

$$R_{B\zeta,\Delta}(t) = \begin{cases} 0 & t < 0 \\ \mathcal{E}^{\scriptscriptstyle T}R_{\Delta^2,\zeta}(0) \quad + O(\varepsilon^3) & t = 0 \\ \mathcal{E}^{\scriptscriptstyle T}A^t R_{\Delta^2,\zeta}(-t) + \mathcal{E}^{\scriptscriptstyle T}\sum_{i=0}^{t-1} A^{t-1-i}\mathcal{E}^{\scriptscriptstyle T}A^i R_\zeta(t-i)\Sigma^{\Delta^\bullet} \quad + O(\varepsilon^3) & t \geq 1 \end{cases} \tag{26}$$

where $A = P_0^{\scriptscriptstyle T}$ *and*

$$R_{\Delta^2,\zeta}(t) = \mathsf{E}[\Delta_t^{\scriptscriptstyle T}\Delta_t\zeta_0] \tag{27}$$

(d) *The cross-covariance* $R_{V\zeta^2,\Delta}(t) = \mathsf{E}[V_t\zeta_t^2\Delta_0]$ *admits the approximation,*

$$R_{V\zeta^2,\Delta}(t) = \begin{cases} 0 & t < 0 \\ \frac{1}{2}\sigma_\zeta^2(P_0^t\mathcal{W})^{\scriptscriptstyle T}\Sigma^{\Delta^\bullet} + O(\varepsilon^3) & t \geq 0 \end{cases} \tag{28}$$

□

The derivation of Theorem 2.4 is given in Appendix E.

Theorem 2.4 leaves out an approximation for π_ε that is required in (23). It also leaves out an approximation for $R_{\Gamma,\zeta}(0)$ required in (23), and approximations for $\{R_{\Delta^2,\zeta}(t) : t \leq 0\}$ in (26). These are obtained in the following:

Proposition 2.5. *The following hold under Assumptions* ***A1*** *and* ***A2****:*

(i) *The steady state mean admits the approximation,*

$$\pi_\varepsilon = \pi_0 + \xi U_1 + O(\varepsilon^3) \tag{29}$$

where U_1 *is the fundamental matrix* (18), *and*

$$\xi = \big(R_{\Gamma,\zeta}(0)\big)^{\scriptscriptstyle T}\mathcal{E} + \tfrac{1}{2}\sigma_\zeta^2\pi_0\mathcal{W}. \tag{30}$$

(ii) *For* $t \geq 0$ *we have,*

$$\begin{aligned} R_{\Delta^2,\zeta}(-t) = {} & \operatorname{diag}(R_{\Gamma,\zeta}(-t-1)P_0) \\ & - P_0^{\scriptscriptstyle T}\operatorname{diag}(R_{\Gamma,\zeta}(-t-1))P_0 + R_\zeta(t+1)\mathsf{E}[\mathcal{X}_\bullet^{(1)}] + O(\varepsilon^3) \end{aligned} \tag{31}$$

where $\mathsf{E}[\mathcal{X}_\bullet^{(1)}] := \operatorname{diag}(\pi_0\mathcal{E}) - (P_0^{\scriptscriptstyle T}\Pi_0\mathcal{E} + [P_0^{\scriptscriptstyle T}\Pi_0\mathcal{E}]^{\scriptscriptstyle T})$.

(iii) *The correlation* $R_{\Gamma,\zeta}$ *is approximated as the infinite sum,*

$$R_{\Gamma,\zeta}(t) = \varepsilon^2\sum_{i=1}^{\infty}\big(B^{\scriptscriptstyle T}P_0^{i-1}\big)^{\scriptscriptstyle T}R_{\zeta^1}(t-i) + O(\varepsilon^3), \qquad t \in \mathbb{Z}\,, \tag{32}$$

in which $\|B^{\scriptscriptstyle T}P_0^{i-1}\| \to 0$ *geometrically fast as* $i \to \infty$.

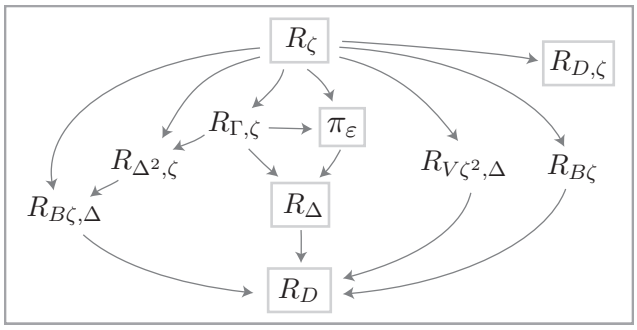

Figure 2: Dependency of autocorrelation functions involved in the approximations of $R(t)$ in (13).

□

The proof of (29) is given in Appendix C, (31) is given in Appendix D.2, and (32) is established in Appendix D.4. The geometric bound on the limit $\|B^T P_0^{i-1}\| \to 0$ follows from Lemma E.2, and the formula $B^T = \pi_0 \mathcal{E}$.

The directed graph shown in Fig. 2 summarizes the dependency between all of these terms. For example, the approximation of $R_{D,\zeta}$ only requires R_ζ, and the covariance Σ^Δ that defines R_Δ is a function of $R_{\Gamma,\zeta}$ and π_ε. The approximation of R_D is a function of the four correlation functions shown (as can also be seen from (21a)–(21d)). The five boxed terms are those that are of interest to us directly; the remaining five terms are introduced only to obtain a closed set of algebraic equations.

Closed-form expressions for the approximations in Proposition 2.5 are possible under **A3**. The proof of Proposition 2.6 is given in Appendix D.4.

Proposition 2.6. *Under **A1**–**A3**, the row vector ξ in (30) has the approximation,*

$$\xi = \varepsilon^2 B^T \sum_{k=1}^{n_z} a_k \rho_k [I - \rho_k P_0]^{-1} \mathcal{E} + \tfrac{1}{2}\varepsilon^2 \sigma_{\zeta^1}^2 \pi_0 \mathcal{W} + O(\varepsilon^3). \tag{33}$$

□

2.3 Power spectral density approximations

Theorem 2.4 provides a second-order approximation of the auto-covariance function $\{R(t)\}$ defined in (13), which we denote $\{\widehat{R}(t)\}$. In particular, $\widehat{R}_D(t)$ is defined as the sum of (21a)–(21d). Based on this and Proposition 2.2 we obtain a second-order approximation $\{\widehat{R}^{\text{tot}}(t)\}$ of the auto-covariance function $\{R^{\text{tot}}(t)\}$ for the triple $(\boldsymbol{\Gamma}, \boldsymbol{D}, \boldsymbol{\zeta})$.

The power spectral density (PSD) of a stationary process is the Fourier transform of its *auto-covariance*. This matrix-valued function is denoted

$$\mathrm{S}(\theta) = \sum_{t=-\infty}^{\infty} \Sigma^{\text{tot}}(t) e^{-j\theta t}, \quad \theta \in \mathbb{R}$$

in which $\Sigma^{\text{tot}}(t) := R^{\text{tot}}(t) - \mu\mu^T$, $t \geq 0$, with $\mu^T = \mathsf{E}[(\Gamma_t, D_t, \zeta_t)]$.

To define an approximation for S we must obtain an approximation $\{\widehat{\Sigma}(t)\}$ that is summable. It turns out that this is obtained from $\{\widehat{R}(t)\}$ without normalization. For each t, the $(2d+1)\times(2d+1)$ matrix is decomposd as follows:

$$\widehat{\Sigma}^{\text{tot}}(t) = \begin{bmatrix} \widehat{\Sigma}_\Gamma(t) & \widehat{\Sigma}_{\Gamma,(D,\zeta)}(t) \\ \widehat{\Sigma}_{(D,\zeta),\Gamma}(t) & \widehat{\Sigma}(t) \end{bmatrix}$$

in which $\widehat{\Sigma}(t) = \widehat{R}(t)$, and the remaining terms are what would be obtained by ignoring the $O(\varepsilon^3)$ error term appearing in (11), and replacing P_0 by its deviation $P_0 - 1\otimes\pi_0$. Denote $\widetilde{A} = (P_0 - 1\otimes\pi_0)^{\scriptscriptstyle T}$, and

$$\widehat{\Sigma}_\Gamma(t) = \sum_{i,j=0}^{\infty} \widetilde{A}^i \widehat{R}_D(t-i+j)(\widetilde{A}^j)^{\scriptscriptstyle T}$$

The matrix $\widehat{\Sigma}_{\Gamma,(D,\zeta)}(t)$ is the $(d+1)$-dimensional column vector whose first d components are

$$\widehat{\Sigma}_{\Gamma,D}(t) = \sum_{i=0}^{\infty} \widetilde{A}^i \widehat{R}_D(t-i)$$

and the final component is defined by the right hand side of (32), ignoring the approximation error. This can be equivalently expressed,

$$\widehat{\Sigma}_{\Gamma,\zeta}(t) = \varepsilon^2 \sum_{i=1}^{\infty} \widetilde{A}^{i-1} B R_{\zeta^1}(t-i)$$

where we have used the fact that $\widetilde{A}^k B = (P_0^{\scriptscriptstyle T})^k B$ since $1^{\scriptscriptstyle T} B = \sum B_i = 0$. Finally, $\widehat{\Sigma}_{(D,\zeta),\Gamma}(t) = \widehat{\Sigma}_{\Gamma,(D,\zeta)}(-t)^{\scriptscriptstyle T}$.

Denote

$$\widehat{\mathrm{S}}(\theta) = \sum_{t=-\infty}^{\infty} \widehat{\Sigma}^{\mathrm{tot}}(t) e^{-j\theta t}, \quad \theta \in \mathbb{R}$$

It can be shown that the sequence $\{\widehat{R}^{\mathrm{tot}}(t)\}$ is absolutely summable, so that the approximation $\widehat{\mathrm{S}}$ is a continuous bounded function of θ. The following is an immediate corollary to Theorem 2.4:

Proposition 2.7. *The approximation of the power spectral density of the stationary sequence* $\{D_{t+1} = \Delta_{t+1} + B_t^{\scriptscriptstyle T}\zeta_t + V_t^{\scriptscriptstyle T}\zeta_t^2\}$ *can be expressed as the sum,*

$$\begin{aligned}\widehat{\mathrm{S}}_D(\theta) = \widehat{\mathrm{S}}_\Delta(\theta) + \widehat{\mathrm{S}}_{B\zeta}(\theta) + e^{-j\theta}\widehat{\mathrm{S}}_{B\zeta,\Delta}(\theta) + \big[e^{-j\theta}\widehat{\mathrm{S}}_{B\zeta,\Delta}(\theta)\big]^* \\ + e^{-j\theta}\widehat{\mathrm{S}}_{V\zeta^2,\Delta}(\theta) + \big[e^{-j\theta}\widehat{\mathrm{S}}_{V\zeta^2,\Delta}(\theta)\big]^*\end{aligned} \tag{34}$$

in which each approximation on the right hand side is obtained from is obtained through the Fourier transform of the corresponding approximations (22)–(28) in Theorem 2.4, and where "" denotes complex-conjugate transpose.*

The power spectral density approximation for $\mathbf{\Gamma}$ *is given by,*

$$\widehat{\mathrm{S}}_\Gamma(\theta) = [Ie^{-j\theta} - \widetilde{A}]^{-1}\widehat{\mathrm{S}}_D(\theta)[Ie^{j\theta} - \widetilde{A}^{\scriptscriptstyle T}]^{-1}$$

and the cross-power spectral density approximations are

$$\widehat{\mathrm{S}}_{\Gamma,D}(\theta) = [Ie^{-j\theta} - \widetilde{A}]^{-1}\widehat{\mathrm{S}}_D(\theta)$$

$$\widehat{\mathrm{S}}_{\Gamma,\zeta}(\theta) = \varepsilon^2 \sum_{t=-\infty}^{\infty}\sum_{i=1}^{\infty} \widetilde{A}^{i-1} B R_{\zeta^1}(t-i)e^{-j\theta t}$$

□

Under slightly stronger assumptions we obtain a uniform bound for this approximation. The proof of Proposition 2.8 is given in Section E.4.

Proposition 2.8. *Suppose that Assumptions* ***A1*** *and* ***A2*** *hold, and in addition* $R_\zeta(t) \to 0$ *geometrically fast as* $t \to \infty$*. Then, the uniform approximation holds: For any* $\varrho \in (0,1)$*,*

$$\widehat{\mathrm{S}}(\theta) = \mathrm{S}(\theta) + O(\varepsilon^{2+\varrho}), \quad \theta \in \mathbb{R}$$

□

3 Example: bits through queues

The following example is motivated by the communication model of [6]. There is a sender that wishes to send data to a receiver. Neither has access to a communication channel in the usual sense. Instead, the sender manipulates the timing of packets to a queue, and the receiver gathers data through observations of the timing of departures from the queue.

3.1 Timing channel model

To obtain a finite state-space model it is assumed that the queue size is bounded by $\bar{q}$, and arrivals are rejected if they cause an overflow. The dynamics of the queue are described as a reflected random walk,

$$Q_{t+1} = \min\{\bar{q}, \max(0, Q_t - S_{t+1} + A_{t+1})\}, \qquad t \geq 0 \tag{35}$$

In the nominal model in which $\boldsymbol{\zeta} \equiv 0$, the pair process $(\boldsymbol{S}, \boldsymbol{A})$ is i.i.d. on $\mathbb{Z}_+^2$. The sender wishes to manipulate the arrival process $\boldsymbol{A}$, and the receiver observes the departure process $\boldsymbol{S}$. This manipulation is modeled through a scalar input sequence $\boldsymbol{\zeta}$.

For simplicity, for the nominal model we restrict to the M/M/1 queue: The usual model evolves in continuous time, but after sampling using *uniformization* one obtains (35), in which $\boldsymbol{A}$ a Bernoulli sequence, and $S_t = 1 - A_t$ for each t. For each integer $n \in \mathsf{X} = \{0, 1, \ldots, \bar{q}\}$, denote $n^+ = \min(n+1, \bar{q})$ and $n^- = \max(n-1, 0)$. If $0 < \lambda < \frac{1}{2}$ is the probably of success for $\boldsymbol{A}$, we then have,

$$\mathsf{P}\{Q(t+1) = n^+ \mid Q(t) = n\} = 1 - \mathsf{P}\{Q(t+1) = n^- \mid Q(t) = n\} = \lambda$$

Its steady-state pmf is given by

$$\pi_0^Q(n) = \kappa \rho^n$$

where $\rho = \lambda/(1-\lambda)$, and $\kappa > 0$ is a normalizing factor.

Recall that the receiver observes departures from the queue, which is equivalent to observations of the sequence $\boldsymbol{S}$. To estimate joint statistics we expand the state space to $X(t) = (Q(t), S(t))$, which evolves on the state space $\mathsf{X} = \{0, 1, \ldots, \bar{q}\} \times \{0, 1\}$. The nominal transition matrix is defined as follows,

$$\begin{aligned} P_0((n,s),(n^+,0)) &= \lambda \\ P_0((n,s),(n^-,1)) &= 1 - \lambda \end{aligned}$$

The first identity holds because a transition from (n, s) to $(n^+, 0)$ means that $A_{t+1} = 1$, in which case $S_{t+1} = 1 - A_{t+1} = 0$. The justification for the second identity is symmetrical. The transition matrix is sparse: $P_0(x, x') = 0$ for all but at most two values of x', regardless of x.

The sender wishes to the manipulate timing of arrivals, which motivates the following formulation for the controlled transition matrix:

$$\begin{aligned} P_\zeta((n,s),(n^+,0)) &= \lambda(1+\zeta) \\ P_\zeta((n,s),(n^-,1)) &= 1 - \lambda(1+\zeta) \end{aligned}$$

in which ζ is constrained to the interval $[-1, 1]$. The state process evolves as the nonlinear state space model (1), with

$$\begin{aligned} Q_{t+1} &= \min\{\bar{q}, \max(0, Q_t - 1 + 2\mathbb{I}\{N_{t+1} \leq \lambda(1+\zeta_t)\}\} \\ S_{t+1} &= \mathbb{I}\{N_{t+1} > \lambda(1+\zeta_t)\}, \qquad t \geq 0 \end{aligned}$$

in which $\boldsymbol{N}$ is i.i.d., with marginal equal to the uniform distribution on $[0, 1]$.

Fig. 3 shows a comparison of the steady-state mean queue length as a function of ε^2 for a numerical example. The linear approximation is obtained from the approximation of π_ε given in Prop. 2.5. Other statistics are shown in Fig. 4. Details can be found in Section 3.3.

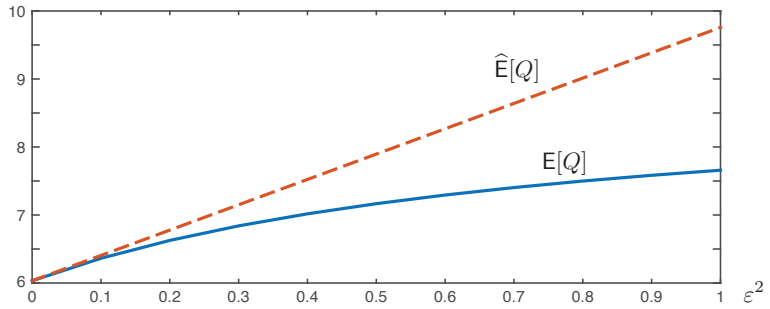

Figure 3: Mean queue length grows approximately linearly in ε^2, for $\varepsilon^2 \le 0.2$. The error for $\varepsilon = 1$ is approximately 25%.

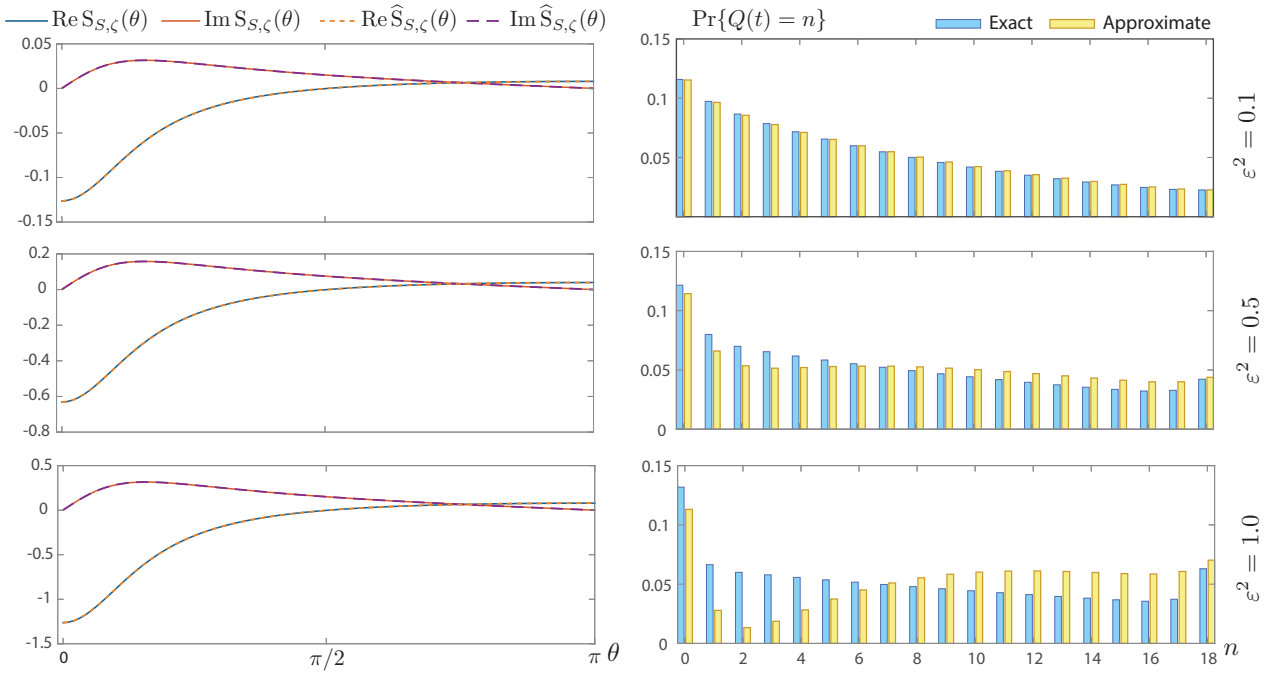

Figure 4: Experiments using $\gamma = 0.4$. The approximation for the cross power spectral density $\mathrm{S}_{S,\zeta}$ appears to be exact for the entire range of ε. The approximation for the steady-state distribution of $\boldsymbol{Q}$ is accurate for $\varepsilon^2 \le 0.5$, but is very poor for $\varepsilon = 1$.

3.2 Second-order bound for mutual information

The mutual information rate $I(\boldsymbol{S}, \boldsymbol{\zeta})$ between $\boldsymbol{S}$ and $\boldsymbol{\zeta}$ defines the capacity of this channel. Letting χ^N denote the joint distribution of $\{S_1, \dots, S_N, \zeta_1, \dots, \zeta_N\}$, and denoting the marginals $\{S_1, \dots, S_N\} \sim \chi_S^N$, $\{\zeta_1, \dots, \zeta_N\} \sim \chi_\zeta^N$, the mutual information rate is defined as the limit

$$I(\boldsymbol{S}, \boldsymbol{\zeta}) = \lim_{N \to \infty} \frac{1}{N} \mathrm{D}(\chi^N \,\|\, \chi_S^N \times \chi_\zeta^N)$$

where D denotes relative entropy (i.e., K-L divergence) [8]. In the following an approximation $\widehat{\mathrm{D}}$ is introduced, and based on this an approximation to mutual information,

$$\hat{I}(\boldsymbol{S}, \boldsymbol{\zeta}) = \lim_{N \to \infty} \frac{1}{N} \widehat{\mathrm{D}}(\chi^N \,\|\, \chi_S^N \times \chi_\zeta^N). \tag{36}$$

The approximation of relative entropy is given here in a general setting. Let ψ_a and ψ_b be probability measures on an abstract measurable space $(\mathcal{E}, \mathcal{B})$. For a measurable function $f \colon \mathcal{E} \to \mathbb{R}$

we let $\psi_a(f)$ denote the mean $\int f(x)\,\psi_a(dx)$. The approximation is the non-negative functional defined as follows:

$$\widehat{\mathrm{D}}(\psi_a\|\psi_b) = \tfrac{1}{2}\sup\Big\{\frac{\psi_a(f)^2}{\psi_b(f^2)} : \ \psi_b(f)=0,\ \psi_a(|f|)<\infty, \quad \textit{and} \quad 0<\psi_b(f^2)<\infty\Big\} \tag{37}$$

The proof of Prop. 3.1 is contained in the Appendix.

Proposition 3.1. *The following hold for any two probability measures ψ_a and ψ_b on an abstract measurable space $(\mathcal{E},\mathcal{B})$. Let $f^* = \log(d\psi_a/d\psi_b)$ denote the log-likelihood ratio.*

(i) *The maximum in (37) is given by $\hat{f}^* = e^{f^*} - 1 = d\psi_a/d\psi_b - 1$, whenever $\widehat{\mathrm{D}}(\psi_a\|\psi_b)$ is finite.*

(ii) *There is a convex, increasing function $\kappa\colon \mathbb{R}_+ \to \mathbb{R}_+$ that vanishes only at the origin, and such that the following bound holds for any two probability measures with bounded log-likelihood ratio:*

$$|\widehat{\mathrm{D}}(\psi_a\|\psi_b) - \mathrm{D}(\psi_a\|\psi_b)| \le \kappa(\|f^*\|_\infty^3)$$

where $\|f^*\|_\infty$ denotes the supremum norm.

□

Returning to the stochastic process setting, in the context of (36), we have for fixed N the following correspondences:

$$\psi_a = \chi^N, \qquad \psi_b = \chi_S^N \times \chi_\zeta^N$$

Consider for $0 \le n < N$ the function

$$f(S_1,\dots,S_N,\zeta_1,\dots,\zeta_N) = \sum_{k=1}^{N-n} \widetilde{S}_{k+n}\zeta_k$$

in which $\widetilde{S}_t = S_t - \mathsf{E}[S_t]$. This has mean zero under ψ_b, and its mean under ψ_a is,

$$\psi_a(f) = \mathsf{E}_{\chi^N}[f(S_1,\dots,S_N,\zeta_1,\dots,\zeta_N)] = (N-n)\Sigma_{S,\zeta}(n)$$

The second moment is also expressed in terms of autocorrelation functions:

$$\begin{aligned}\psi_b(f^2) &= \mathsf{E}_{\chi_S^N\times\chi_\zeta^N}[f^2(S_1,\dots,S_N,\zeta_1,\dots,\zeta_N)] \\ &= \sum_{k=1}^{N-n}\sum_{\ell=1}^{N-n} \mathsf{E}_{\chi_S^N\times\chi_\zeta^N}\mathsf{E}[\widetilde{S}_{k+n}\widetilde{S}_{\ell+n}\zeta_k\zeta_\ell] \\ &= \sum_{k=1}^{N-n}\sum_{\ell=1}^{N-n} \Sigma_S(k-\ell)\Sigma_\zeta(k-\ell)\end{aligned}$$

For fixed n and $N \gg n$ this admits the approximation $\psi_b(f^2) \approx (N-n)\mathrm{S}_{S\times\zeta}(0)$, where

$$\mathrm{S}_{S\times\zeta}(0) = \sum_{m=-\infty}^{\infty} \Sigma_S(m)\Sigma_\zeta(m)\,.$$

This gives the large-N approximation,

$$\widehat{\mathrm{D}}(\psi_a\|\psi_b) \ge \tfrac{1}{2}\frac{\psi_a(f)^2}{\psi_b(f^2)} = \tfrac{1}{2}\frac{\Sigma_{S,\zeta}(n)^2}{\mathrm{S}_{S\times\zeta}(0)}N + O(1)$$

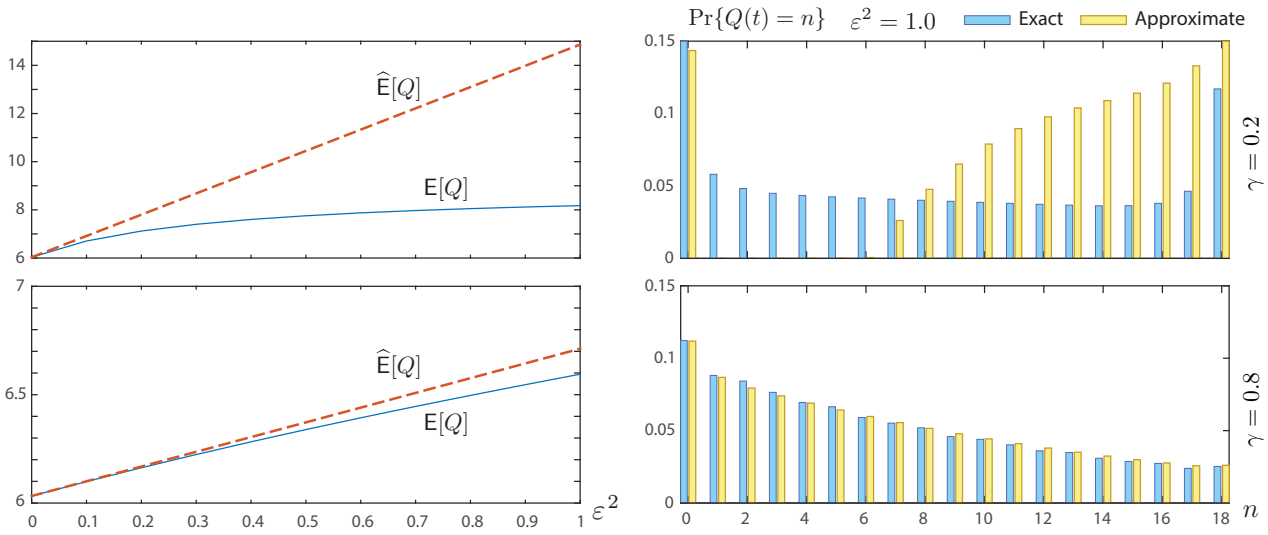

Figure 5: The top row shows results from numerical experiments with $\gamma = 0.2$, and the bottom row $\gamma = 0.8$. The approximation for $\pi_\varepsilon^Q(n)$ is nearly perfect in the latter case, even with $\varepsilon = 1$. With $\gamma = 0.2$, the second order approximation is poor for $\varepsilon \geq 0.5$.

While the derivation was performed for $n \geq 0$, similar arguments establish the same bound for any integer n. The approximation for mutual information rate is thus lower bounded,

$$\hat{I}(\boldsymbol{S}, \boldsymbol{\zeta}) \geq \tfrac{1}{2} \sup_{-\infty<n<\infty} \frac{\Sigma_{S,\zeta}(n)^2}{\mathrm{S}_{S\times\zeta}(0)} \tag{38}$$

This function class is of course highly restrictive. A larger class of functions can be obtained by defining for each n and each $\alpha, \beta \in \mathbb{R}^{n+1}$,

$$f(S_1, \ldots, S_N, \zeta_1, \ldots, \zeta_N) = \sum_{k=1}^{N-n} S_k^\alpha \zeta_k^\beta, \qquad \textit{with} \quad S_k^\alpha = \sum_{m=0}^{n} \alpha_m \widetilde{S}_{k+m}, \quad \zeta_k^\beta = \sum_{m=0}^{n} \beta_m \zeta_{k+m}.$$

Formulae for $\psi_a(f)$ and $\psi_b(f^2)$ can be obtained as in the foregoing, yielding

$$\hat{I}(\boldsymbol{S}, \boldsymbol{\zeta}) \geq \tfrac{1}{2} \frac{\Sigma_{S^\alpha, \zeta^\beta}(0)^2}{\mathrm{S}_{S^\alpha \times \zeta^\beta}(0)}$$

3.3 Numerical experiments

In all of the numerical examples described here, λ is chosen so that $\rho = \lambda/(1-\lambda) = 0.9$, and the upper bound appearing in (35) is $\bar{q} = 18$.

A Markovian model was chosen for ζ^1 so that exact computations can be obtained for the larger Markov chain. A simple model was chosen, in which $\boldsymbol{\zeta}^1$ evolves on the three states $\{-1, 0, 1\}$. The larger Markov chain $\Psi_t = (Q_t, S_t, \zeta_t^1)$ then evolves on a state space of size $6(1 + \bar{q})$.

The three states are labeled $\{z^i : i = 1, 2, 3\} = \{-1, 0, 1\}$. For a fixed parameter $\gamma \in (0, \frac{1}{2})$, the transition matrix K is defined as follows. First, $\mathsf{P}\{\zeta_{t+1} = z^j \mid \zeta_t = z^i\} = \gamma$ whenever $|z^j - z^i| = 1$:

$$K_{1,2} = K_{2,1} = K_{2,3} = K_{3,2} = \gamma$$

The remaining transition probabilities are $K_{1,1} = K_{3,3} = 1-\gamma$, and $K_{2,2} = 1-2\gamma$. The steady-state pmf μ_0 is uniform, so the steady-state variance is

$$\sigma_\zeta^2 = ((-1)^2 + 0^2 + 1^2)/3 = 2/3$$

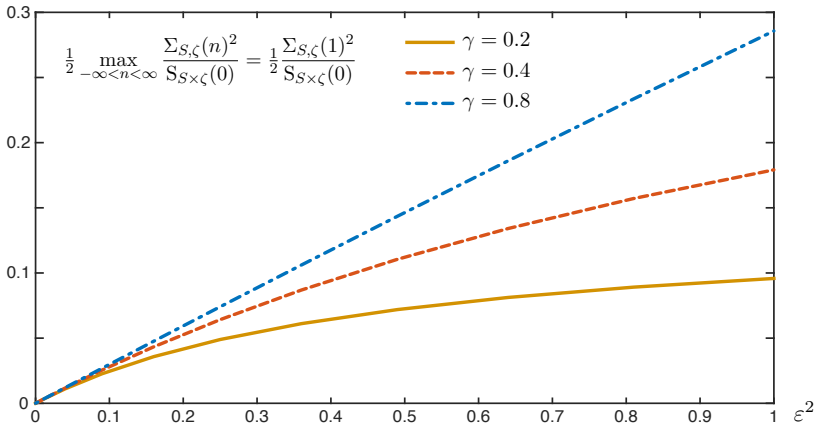

Figure 6: Lower bound for $\hat{I}(\boldsymbol{S},\boldsymbol{\zeta})$ as a function of ε^2 for three values of γ.

Its autocorrelation is equal to its autocovariance: $R_\zeta(m)=\sigma_\zeta^2(1-\gamma)^{|m|}$.

Unless explicitly stated otherwise, the results that follow use $\gamma=0.4$, so that the asymptotic variance (the variance appearing in the Central Limit Theorem) is

$$\text{asym. variance}=\mathrm{S}_\zeta(0)=\sum_{k=-\infty}^{\infty}R_\zeta(k)=\Big(\frac{2}{\gamma}-1\Big)\sigma_\zeta^2=4\sigma_\zeta^2$$

Let $\hat{\pi}_\varepsilon=\pi_0+\xi U_1$ (with ξ and U_1 defined in (2.5)). The approximate pmf illustrated in the plots on the right hand side of Fig. 4 are defined by the first marginal, $\hat{\pi}_\varepsilon^Q(n)=\sum_{s=0,1}\hat{\pi}_\varepsilon(n,s)$, for $n=0,\dots,18$. The approximate steady-state queue length plotted in Fig. 3 is defined by

$$\widehat{\mathsf{E}}[Q]=\sum_{n=0}^{18}n\hat{\pi}_\varepsilon^Q(n)$$

The steady state pmf for $\boldsymbol{\Psi}$ was computed to obtain the exact steady-state mean $\mathsf{E}[Q]$, which is the concave plot shown in Fig. 3. The approximations are accurate for $\varepsilon\le 0.7$.

The approximations for the cross power spectral density shown on the left hand side of Fig. 4 are remarkable.

The statistics of $\boldsymbol{Q}$ and its approximations are highly sensitive to the parameter γ. For $\gamma=0.8$, the approximation of the steady-state mean $\mathsf{E}[Q]$ and approximations of $\pi_\varepsilon^Q(n)$ are nearly exact for the entire range of ε. For $0<\gamma\le 0.2$ the approximations are accurate only for a very narrow range of $\varepsilon\sim 0$. Results for $\gamma=0.2$ and $\gamma=0.8$ are shown in Fig. 5.

The approximation for mutual information in (38) is defined as a maximum, which was achieved at $n=1$ in each experiment:

$$\hat{I}(\boldsymbol{S},\boldsymbol{\zeta})\ge\tfrac{1}{2}\max_{-\infty<n<\infty}\frac{\Sigma_{S,\zeta}(n)^2}{\mathrm{S}_{S\times\zeta}(0)}=\frac{1}{2}\frac{\Sigma_{S,\zeta}(1)^2}{\mathrm{S}_{S\times\zeta}(0)}$$

Plots for four different values of γ are shown in Fig. 6.

The plots in Fig. 6 use the approximations obtained in Section 2. However, plots obtained using the exact values of $\Sigma_{S,\zeta}(1)$ and $\mathrm{S}_{S\times\zeta}(0)$ are indistinguishable.

Unfortunately, we cannot compare $\hat{I}(\boldsymbol{S},\boldsymbol{\zeta})$ with the true mutual information rate. This is a significant computational challenge that is beyond the scope of this paper.

4 Conclusions

It is very surprising to obtain an exact second order Taylor series expansion for these second order statistics with minimal assumptions on the controlled Markov model. The accuracy of the approximations obtained in numerical examples is also fortunate. The companion paper [3] shows that these approximations are also accurate in applications to distributed control. Further work is needed to see if this will lead to useful bounds in applications to information theory.

References

[1] N. V. Krylov, R. S. Lipster, and A. A. Novikov, "Kalman filter for Markov processes," in *Statistics and Control of Stochastic Processes*. New York: Optimization Software, inc., 1984, pp. 197–213.

[2] Y. Chen, A. Bušić, and S. Meyn, "State estimation and mean field control with application to demand dispatch," in *54th Annual Conference on Decision and Control*, Dec 2015, pp. 6548–6555.

[3] ——, "Individual risk in mean field control with application to demand dispatch (in preparation)," *IEEE Trans. on Smart Grid*, 2015.

[4] ——, "Individual risk in mean-field control models for decentralized control, with application to automated demand response," in *Proc. of the 53rd IEEE Conference on Decision and Control*, Dec. 2014, pp. 6425–6432.

[5] P. J. Schweitzer, "Perturbation theory and finite Markov chains," *J. Appl. Prob.*, vol. 5, pp. 401–403, 1968.

[6] V. Anantharam and S. Verdu, "Bits through queues," *IEEE Trans. Inform. Theory*, vol. 42, no. 1, pp. 4–18, 2006.

[7] P. E. Caines, *Linear Stochastic Systems*. New York: John Wiley & Sons, 1988.

[8] T. M. Cover and J. A. Thomas, *Elements of information theory*. New York: John Wiley & Sons Inc., 1991, a Wiley-Interscience Publication.

[9] S. Foss and T. Konstantopoulos, "An overview of some stochastic stability methods," *Journal of the Operations Research Society of Japan-Keiei Kagaku*, vol. 47, no. 4, pp. 275–303, 2004.

[10] A. Dembo and O. Zeitouni, *Large Deviations Techniques And Applications*, 2nd ed. New York: Springer-Verlag, 1998.

Appendices

A Coupling

We present here the proof of Proposition 2.3.

We first obtain a recursion for the joint process $\boldsymbol{\Psi}:=(\boldsymbol{\Gamma},\boldsymbol{\Gamma}^\bullet)$, driven by $\boldsymbol{\zeta}$, and two i.i.d. sequences $\boldsymbol{N}^\circ$ and $\boldsymbol{N}^\bullet$, each with marginal distribution uniform on $[0,1]$. The three sequences $\boldsymbol{\zeta},\boldsymbol{N}^\circ,\boldsymbol{N}^\bullet$ are mutually independent.

Letting $\boldsymbol{W}$ denote the 3-dimensional stationary stochastic process $(\boldsymbol{\zeta},\boldsymbol{N}^\circ,\boldsymbol{N}^\bullet)$, we construct a function F for which,

$$\Psi_{t+1}=F(\Psi_t,W_{t+1}) \tag{39}$$

Since $\boldsymbol{\Psi}$ evolves on a finite set, the existence of a stationary solution follows from [9] (see Theorem 5 and the discussion that follows).

To construct the function F it is enough to define the matrix sequence $\boldsymbol{G}$ that appears in (5), and also the sequence $\boldsymbol{G}^\bullet$ that defines the dynamics of $\boldsymbol{\Gamma}^\bullet$. Each are based on the following definition: for $\zeta\in\mathbb{R}$ and $s\in[0,1]$, denote

$$\mathcal{G}_{i,j}(\zeta,s)=\mathbb{I}\Big\{\sum_{k=1}^{j-1}P_\zeta(i,k)\le s<\sum_{k=1}^{j}P_\zeta(i,k)\Big\}$$

with the convention that '"$\sum_{k=1}^{0}[\cdot]_k=0$". We then take, for any t,

$$G_t^\bullet=\mathcal{G}(0,N_t^\bullet)$$

A third i.i.d. sequence $\boldsymbol{N}$ is obtained by sampling:

$$N_{t+1}=\begin{cases}N_{t+1}^\bullet & \text{if } \Gamma_t=\Gamma_t^\bullet\\ N_{t+1}^\circ & \text{else}\end{cases} \tag{40}$$

We then take $G_t=\mathcal{G}(\zeta_{t-1},N_t)$.

Based on these definitions, the evolution equation (39) holds for some F; we now focus analysis on a stationary solution defined for all $t\in\mathbb{Z}$.

Choose $T_0\ge1$, $\delta_0>0$, $\varepsilon_0>0$, so that

$$\begin{aligned}\mathsf{P}\{X_{t+1}=X_{t+1}^\bullet\mid X_t=X_t^\bullet\}&\ge1-\delta_0\varepsilon\\ \mathsf{P}\{X_{t+T_0}=X_{t+T_0}^\bullet\mid X_t\ne X_t^\bullet\}&\ge\delta_0,\qquad\qquad t\in\mathbb{Z},\ \varepsilon\le\varepsilon_0\end{aligned}$$

This is possible by the construction of the joint evolution equations, and the assumption that $\boldsymbol{X}^\bullet$ is irreducible and aperiodic. The first bound may be extended to obtain,

$$\mathsf{P}\{X_{t+T}=X_{t+T}^\bullet\mid X_t=X_t^\bullet\}\ge(1-\delta_0\varepsilon)^T\ge1-T\delta_0\varepsilon,\quad T\ge1.$$

We then have by stationarity,

$$\begin{aligned}\mathsf{P}\{X_0=X_0^\bullet\}=\mathsf{P}\{X_{T_0}=X_{T_0}^\bullet\}&=\mathsf{P}\{X_{T_0}=X_{T_0}^\bullet\mid X_0=X_0^\bullet\}\mathsf{P}\{X_0=X_0^\bullet\}\\&\quad+\mathsf{P}\{X_{T_0}=X_{T_0}^\bullet\mid X_0\ne X_0^\bullet\}\mathsf{P}\{X_0\ne X_0^\bullet\}\end{aligned}$$

Now, substitute the prior bounds, giving

$$\begin{aligned}\mathsf{P}\{X_0=X_0^\bullet\}&\ge(1-T_0\delta_0\varepsilon)\mathsf{P}\{X_0=X_0^\bullet\}\\&\quad+\delta_0\mathsf{P}\{X_0\ne X_0^\bullet\}\end{aligned}$$

Substituting $\mathsf{P}\{X_0\ne X_0^\bullet\}=1-\mathsf{P}\{X_0=X_0^\bullet\}$ and rearranging terms gives,

$$\mathsf{P}\{X_0=X_0^\bullet\}\ge\frac{1}{1+T_0\varepsilon}\ge1-T_0\varepsilon$$

which completes the proof of (14).

The approximation (15) follows from (14) and independence of $\boldsymbol{\zeta}$ and $\boldsymbol{\Gamma}^\bullet$:

$$\mathsf{E}[\Gamma_t\zeta_t]=\mathsf{E}[\zeta_t\Gamma_t^\bullet]+\mathsf{E}[\zeta_t(\Gamma_t-\Gamma_t^\bullet)]=\mathsf{E}[\zeta_t]\mathsf{E}[\Gamma_t^\bullet]+O(\varepsilon^2)=O(\varepsilon^2)$$

The remaining bounds in (16) follow directly from (14) and the smoothnes assumptions on P_ζ. □

B Martingale difference sequence

Here we give the proof of Proposition 2.1. Define the σ-algebra $\mathcal{F}_t = \sigma\{\zeta_{-\infty}^{\infty}, N_{-\infty}^{t}\}$. The random vector Δ_t is $\mathcal{F}_t$-measurable for each t, and it follows from (6) that

$$\mathsf{E}[\Delta_{t+1}|\mathcal{F}_t] = \mathsf{E}[\Gamma_t[\mathcal{G}(\zeta_t, N_{t+1}) - P_{\zeta_t}]|\mathcal{F}_t] = 0$$

This proves that Δ_t is a martingale difference sequence. Moreover, using the smoothing property of conditional expectation, for any t and τ,

$$\begin{aligned}\mathsf{E}[\Delta_{t+1}\zeta_\tau] &= \mathsf{E}[\Gamma_t[\mathcal{G}(\zeta_t, N_{t+1}) - P_{\zeta_t}]\zeta_\tau] \\ &= \mathsf{E}[\Gamma_t\zeta_\tau\ \mathsf{E}[\{\mathcal{G}(\zeta_t, N_{t+1}) - P_{\zeta_t}\}|\mathcal{F}_t]] = 0.\end{aligned}$$

This establishes $R_{\Delta,\zeta}(k) = 0$ for any k, which is (10).

To obtain the covariance of $\mathbf{\Delta}$, we first compute the second moment of $\Gamma_t G_{t+1}$. We use the fact that Γ_t and G_{t+1} have entries zero or one in the following:

$$\begin{aligned}\mathsf{E}[(\Gamma_t G_{t+1})^{\mathrm{T}}\Gamma_t G_{t+1}]_{ij} &= \sum_{k=1}^{d}\mathsf{E}[\Gamma_t(x^k)G_{t+1}(k,i)G_{t+1}(k,j)] \\ &= \sum_{k=1}^{d}\mathsf{E}[\Gamma_t(x^k)G_{t+1}(k,i)]\mathbb{I}\{i=j\}\end{aligned}$$

We also have $\mathsf{E}[\Gamma_t G_{t+1}] = \mathsf{E}[\Gamma_t P_{\zeta_t}]$. These two identities imply the formula,

$$\mathsf{E}[(\Gamma_t G_{t+1})^{\mathrm{T}}\Gamma_t G_{t+1}] = \operatorname{diag}(\mathsf{E}[\Gamma_t P_{\zeta_t}]) = \operatorname{diag}(\mathsf{E}[\Gamma_{t+1}]).$$

The result (9) follows from this second moment formula:

$$\begin{aligned}\operatorname{Cov}(\Gamma_t[G_{t+1} - P_{\zeta_t}]) &= \mathsf{E}[(\Gamma_t G_{t+1})^{\mathrm{T}}\Gamma_t G_{t+1}] - \mathsf{E}[(\Gamma_t P_{\zeta_t})^{\mathrm{T}}\Gamma_t P_{\zeta_t}] \\ &= \operatorname{diag}(\mathsf{E}[\Gamma_{t+1}]) - \mathsf{E}[(P_{\zeta_t})^{\mathrm{T}}\operatorname{diag}(\Gamma_t)P_{\zeta_t}].\end{aligned}$$

□

C Approximating the steady-state mean

The approximation (29) is obtained here, starting with the approximate evolution equation that was used to obtain (11):

$$\begin{aligned}\Gamma_{t+1} &= \Gamma_t P_{\zeta_t} + \Delta_{t+1} \\ &= \Gamma_t[P_0 + \zeta_t\mathcal{E} + \tfrac{1}{2}\zeta_t^2\mathcal{W}] + \Delta_{t+1} + O(\varepsilon^3)\end{aligned}$$

Taking the mean of each side, and using stationarity,

$$\begin{aligned}\mathsf{E}[\Gamma_t] = \mathsf{E}[\Gamma_{t+1}] = \mathsf{E}[\Gamma_t]P_0 \\ + \mathsf{E}[\zeta_t\Gamma_t]\mathcal{E} \\ + \tfrac{1}{2}\mathsf{E}[\zeta_t^2\Gamma_t]\mathcal{W} + O(\varepsilon^3)\end{aligned} \tag{41}$$

To approximate $\mathsf{E}[\zeta_t^2\Gamma_t]$ we use $\Gamma_t = \Gamma_t^\bullet + \widetilde{O}(\varepsilon)$. This combined with independence of $\mathbf{\Gamma}^\bullet, \boldsymbol{\zeta}$ gives,

$$\mathsf{E}[\zeta_t^2\Gamma_t] = \mathsf{E}[\zeta_t^2\Gamma_t^\bullet] + \mathsf{E}[\zeta_t^2(\Gamma_t - \Gamma_t^\bullet)] = \mathsf{E}[\zeta_t^2]\mathsf{E}[\Gamma_t^\bullet] + O(\varepsilon^3) = \sigma_\zeta^2\pi_0 + O(\varepsilon^3)$$

Substituting this into (41) gives the approximate fixed-point equation,

$$\tilde{\pi}_\varepsilon = \tilde{\pi}_\varepsilon P_0 + \xi + O(\varepsilon^3) \tag{42}$$

where $\tilde{\pi}_\varepsilon = \pi_\varepsilon - \pi_0$, and ξ is defined in (30).

The matrix $I - P_0$ is not invertible since it has eigenvalue at 0. To obtain an invertible matrix we note that $\tilde{\pi}_\varepsilon[1 \otimes \pi_0] = 0$ for any ε, and hence (42) is equivalent to the vector equation,

$$\tilde{\pi}_\varepsilon[I - P_0 + 1 \otimes \pi_0] = \tilde{\pi}_\varepsilon[I - P_0] = \xi\varepsilon^2 + O(\varepsilon^3)$$

The desired result (29) is obtained by inversion:

$$\tilde{\pi}_\varepsilon = \xi[I - P_0 + 1 \otimes \pi_0]^{-1} + O(\varepsilon^3) = \xi U_1 + O(\varepsilon^3)$$

□

D Cross-covariance with ζ

Approximations for $R_{D,\zeta}$ and $R_{B\zeta}$ are relatively simple because $\zeta_t = O(\varepsilon)$.

D.1 Cross-covariance between D and ζ

Using the coupling result we obtain an approximation for the cross-correlation function,

$$\begin{aligned} R_{D,\zeta}(t) &= \mathsf{E}[D_{t+1}^T \zeta_1] \\ &= \mathsf{E}[(B_t^T \zeta_t + \Delta_{t+1})^T \zeta_1] + O(\varepsilon^3) \\ &= \mathsf{E}[B_t^\bullet \zeta_t \zeta_1] + O(\varepsilon^3) \\ &= B R_\zeta(t-1) + O(\varepsilon^3). \end{aligned} \tag{43}$$

D.2 Cross-covariance between $\Delta\Delta^T$ and ζ

Recall the σ-algebra $\mathcal{F}_s = \sigma\{\zeta_{-\infty}^\infty, N_{-\infty}^s\}$ introduced in Section B. Taking $s = -1$ we obtain from the smoothing property of the conditional expectation,

$$R_{\Delta^2,\zeta}(-t) = \mathsf{E}[\Delta_0^T \Delta_0 \zeta_t] = \mathsf{E}[\zeta_t \mathsf{E}[\Delta_0^T \Delta_0 \mid \mathcal{F}_{-1}]], \quad t \in \mathbb{Z}.$$

The conditional expectation is a matrix that is denoted

$$\mathcal{X} = \mathsf{E}[\Delta_0^T \Delta_0 | \mathcal{F}_{-1}] = \operatorname{diag}(\Gamma_{-1} P_{\zeta_{-1}}) - P_{\zeta_{-1}}^T \Lambda_{-1}^\Gamma P_{\zeta_{-1}} \tag{44}$$

Lemma D.1. *For $t \geq 0$ we have*

$$\begin{aligned} R_{\Delta^2,\zeta}(-t) &= \mathsf{E}[\zeta_t \mathcal{X}] \\ &= \operatorname{diag}(R_{\Gamma,\zeta}(-t-1)P_0) - P_0^T \operatorname{diag}(R_{\Gamma,\zeta}(-t-1))P_0 + R_\zeta(t+1)\mathsf{E}[\mathcal{X}_\bullet^{(1)}] + O(\varepsilon^3), \end{aligned}$$

where $\mathsf{E}[\mathcal{X}_\bullet^{(1)}]$ is defined below (31).

Proof. A first order Taylor series approximation gives $\mathcal{X} = \mathcal{X}^{(0)} + \zeta_{-1}\mathcal{X}^{(1)} + O(\varepsilon^2)$, where

$$\begin{aligned}\mathcal{X}^{(0)} &= \operatorname{diag}(\Gamma_{-1}P_0) - P_0^T \operatorname{diag}(\Gamma_{-1})P_0 \\ \mathcal{X}^{(1)} &= \operatorname{diag}(\Gamma_{-1}\mathcal{E}) - (P_0^T \operatorname{diag}(\Gamma_{-1})\mathcal{E} + [P_0^T \operatorname{diag}(\Gamma_{-1})\mathcal{E}]^T)\,.\end{aligned}$$

Hence, for $t \ge 0$,

$$\begin{aligned}\mathsf{E}[\zeta_t\mathcal{X}] &= \mathsf{E}[\zeta_t(\mathcal{X}^{(0)} + \zeta_{-1}\mathcal{X}^{(1)} + O(\varepsilon^2))] \\ &= \mathsf{E}[\zeta_t\mathcal{X}^{(0)}] + \mathsf{E}[\zeta_t\zeta_{-1}\mathcal{X}_\bullet^{(1)}] + O(\varepsilon^3) \\ &= \mathsf{E}[\zeta_t\mathcal{X}^{(0)}] + R_\zeta(t+1)\mathsf{E}[\mathcal{X}_\bullet^{(1)}] + O(\varepsilon^3)\end{aligned}$$

where, in the second equality we used $\mathcal{X}^{(1)} = \mathcal{X}_\bullet^{(1)} + \widetilde{O}(\varepsilon)$ with

$$\mathcal{X}_\bullet^{(1)} = \operatorname{diag}(\Gamma_{-1}^\bullet\mathcal{E}) - \left(P_0^T \operatorname{diag}(\Gamma_{-1}^\bullet)\mathcal{E} + [P_0^T \operatorname{diag}(\Gamma_{-1}^\bullet)\mathcal{E}]^T\right)\,,$$

and also used the fact that $\Gamma_{-1}^\bullet$ is independent of $\boldsymbol{\zeta}$. We have by the definitions,

$$\mathsf{E}[\zeta_t\mathcal{X}^{(0)}] = \operatorname{diag}(R_{\Gamma,\zeta}(-t-1)P_0) - P_0^T \operatorname{diag}(R_{\Gamma,\zeta}(-t-1))P_0.$$

Substitution into the previous approximation for $\mathsf{E}[\zeta_t\mathcal{X}]$ completes the proof. □

D.3 Auto-correlation of $B\boldsymbol{\zeta}$

Applying the coupling result (16) in Proposition 2.3 gives $B_t^T = B^\bullet{}_t^T + \widetilde{O}(\varepsilon)$, where $B_t^T = \Gamma_t\mathcal{E}$. Hence,

$$R_{B\zeta}(t) = \mathsf{E}[B_t\zeta_t B_0^T\zeta_0] = \mathsf{E}[B_t^\bullet\zeta_t(B_0^\bullet\zeta_0)^T] + O(\varepsilon^3) = R_{B^\bullet\zeta}(t) + O(\varepsilon^3) \tag{45}$$

Independence of $\boldsymbol{\zeta}$ and $\Gamma_t^\bullet$ implies a formula for the simpler auto-correlation:

$$\begin{aligned}R_{B^\bullet\zeta}(t) &:= \mathsf{E}[B_t^\bullet\zeta_t(B_0^\bullet\zeta_0)^T] \\ &= \mathsf{E}[(\Gamma_t^\bullet\mathcal{E})^T(\Gamma_t^\bullet\mathcal{E})]\mathsf{E}[\zeta_t\zeta_0] \\ &= \mathcal{E}^T R_{\Gamma^\bullet}(t)\mathcal{E}R_\zeta(t)\end{aligned} \tag{46}$$

A formula for $R_{\Gamma^\bullet}(t)$ is given next: For $t \ge 0$,

$$R_{\Gamma^\bullet}(t) = \mathsf{E}[(\Gamma_t^\bullet)^T\Gamma_0^\bullet] = \mathsf{E}[(\Gamma_0^\bullet P_0^t)^T\Gamma_0^\bullet] = (P_0^T)^t\mathsf{E}[(\Gamma_0^\bullet)^T\Gamma_0^\bullet] = (P_0^T)^t\Pi_0 \tag{47}$$

where the last equality is from the fact that Γ_t has binary entries and $\mathsf{E}[\Gamma_t] = \pi_0$.

Combining (45)–(47) completes the proof of (22). □

D.4 Cross-covariance between Γ and $\boldsymbol{\zeta}$

The derivation of (32) is obtained via a recursion, similar to the calculation in Section C. It is simplest to work with the row vectors $\nu_k = \big(R_{\Gamma,\zeta}(k)\big)^T = \mathsf{E}[\Gamma_k\zeta_0]$. For any $k \in \mathbb{Z}$,

$$\begin{aligned}\nu_{k+1} &= \mathsf{E}[\Gamma_{k+1}\zeta_0] \\ &= \mathsf{E}[(\Gamma_k P_{\zeta_k} + \Delta_{k+1})\zeta_0] \\ &= \mathsf{E}[\Gamma_k P_{\zeta_k}\zeta_0] \\ &= \mathsf{E}[\Gamma_k(P_0 + \mathcal{E}\zeta_k + O(\varepsilon^2))\zeta_0] \\ &= \nu_k P_0 + \mathsf{E}[\Gamma_k\mathcal{E}\zeta_k\zeta_0] + O(\varepsilon^3)\end{aligned}$$

where in the third equation we used the fact that the sequences $\mathbf{\Delta}$ and $\boldsymbol{\zeta}$ are uncorrelated. Recalling the definition $B^T = \pi\mathcal{E} = \mathsf{E}[\Gamma_k^\bullet]\mathcal{E}$, and applying the coupling result $\Gamma_k = \Gamma_k^\bullet + \widetilde{O}(\varepsilon)$ gives

$$\mathsf{E}[\Gamma_k\mathcal{E}\zeta_k\zeta_0] = \mathsf{E}[\Gamma_k^\bullet\mathcal{E}\zeta_k\zeta_0] + O(\varepsilon^3) = B^T R_\zeta(k) + O(\varepsilon^3).$$

Hence, $\nu_{k+1} = \nu_k P_0 + B^T R_\zeta(k) + O(\varepsilon^3)$; that is, there is a bounded sequence of row vectors $\{\gamma_k\}$ such that

$$\nu_{k+1} = \nu_k P_0 + B^T R_\zeta(k) + \varepsilon^3\gamma_k$$

It follows from Lemma E.2 that $B^T 1 = 0$. Moreover, since Γ_{t-n} is a pmf, we have for any k, ℓ,

$$\nu_k 1 = \mathsf{E}[\zeta_k\Gamma_\ell]1 = \mathsf{E}[\zeta_k] = 0. \tag{48}$$

It then follows that $\gamma_k 1 = 0$ for each k.

On iterating, we obtain for any integer $n \geq 1$,

$$\nu_{k+n} = \nu_k P_0^n + \sum_{i=1}^{n}[B^T + \varepsilon^3\gamma_i]P_0^{i-1}R_\zeta(k+n-i).$$

Now, substitute $t = k + n$, where $t \in \mathbb{Z}$ is a fixed integer, and n is a large positive integer:

$$\gamma_t = R_{\Gamma,\zeta}(t-n)P_0^n + \sum_{i=1}^{n}[B^T + \varepsilon^3\gamma_i]P_0^{i-1}R_\zeta(t-i). \tag{49}$$

For large n we have $P_0^n = 1\otimes\pi_0 + o_e(1)$, where $o_e(1) \to 0$ geometrically fast as $n \to \infty$. Consequently,

$$\begin{aligned} R_{\Gamma,\zeta}(t-n)P_0^n &= \mathsf{E}[\zeta_0\Gamma_{t-n}]1 \otimes \pi_0 + o_e(1) \\ \gamma_i P_0^{i-1} &= \gamma_i 1 \otimes \pi_0 + o_e(1) \end{aligned}$$

We have seen previously following (48) that $\gamma_i 1 = 0$. It follows similarly that $\mathsf{E}[\zeta_0\Gamma_{t-n}]1 = \mathsf{E}[\zeta_0] = 0$, from which we conclude that

$$R_{\Gamma,\zeta}(t-n)P_0^n = o_e(1)$$

$$\textit{and} \quad \sum_{i=1}^{\infty}\|\gamma_i P_0^{i-1}R_\zeta(t-i)\| < \infty$$

This justifies letting $n \to \infty$ in (49) to obtain,

$$\gamma_t = B^T\sum_{i=1}^{\infty}P_0^{i-1}R_\zeta(t-i) + O(\varepsilon^3)\,,$$

which is equivalent to (32). □

Based on this result we now establish Proposition 2.6. It is sufficient to establish the following approximation:

$$R_{\Gamma,\zeta}(0) = \varepsilon^2\sum_{k=1}^{n_z}a_k\rho_k[I - \rho_k P_0^T]^{-1}B + O(\varepsilon^3) \tag{50}$$

The representation (33) for ξ then follows immediately from the original definition (30).

Recall that $\nu_k = \mathsf{E}[\Gamma_k\zeta_0]$, $k \in \mathbb{Z}$. Under **A3** we apply (17) to conclude that for $t \leq 1$ and $i \geq 1$,

$$R_\zeta(t-i) = \varepsilon^2\sum_{k=1}^{n_z}a_k\rho_k^{i-t}\,.$$

For these values of t and i we have $i - t = i + |t|$, and hence the approximation (32) gives,

$$\nu_t = \varepsilon^2 \sum_{k=1}^{n_z} \sum_{i=1}^{\infty} B^{\scriptscriptstyle T} P_0^{i-1} a_k \rho_k^{i+|t|} + O(\varepsilon^3).$$

Rearranging terms, and letting $j = i - 1$ gives,

$$\nu_t = \varepsilon^2 B^{\scriptscriptstyle T} \sum_{k=1}^{n_z} a_k \rho_k^{1+|t|} \sum_{j=0}^{\infty} P_0^j \rho_k^j + O(\varepsilon^3)\,.$$

On setting $t = 0$ and taking transposes, this becomes (50) □

E Proof of Theorem 2.4

The representation of $R_D(t)$ as the sum of the terms (21a–21d), plus an error of order ε^3, follows immediately from the arguments preceding the theorem. The proof of (20) is given in (43).

Next we consider each of the terms (21a–21d) separately.

The approximation for $R_{B\zeta}(t)$ that appears in (21a) was given in Section D.3. Consider next the covariance (21b).

E.1 Computation of Σ^{Δ} in (23)

Take the expectation inside the diag operator in (9) to obtain

$$\Sigma^{\Delta} = \Pi_{\varepsilon} - \mathsf{E}[P_{\zeta_t}^{\scriptscriptstyle T} \Lambda_t^{\Gamma} P_{\zeta_t}], \tag{51}$$

where, $\Pi_{\varepsilon} = \operatorname{diag}(\pi_{\varepsilon})$, with π_{ε} estimated in Proposition 2.5.

We approximate $\mathsf{E}[P_{\zeta_t}^{\scriptscriptstyle T} \Lambda_t^{\Gamma} P_{\zeta_t}]$ within the expectation using a second order Taylor series expansion. The random matrix Λ_t^{Γ} has binary entries, so we leave it fixed in this approximation.

For any scalar ζ and matrix Λ we have the Taylor series approximation,

$$P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta} = P_0^{\scriptscriptstyle T} \Lambda P_0 + \zeta M^{(1)} + \tfrac{1}{2}\zeta^2 M^{(2)} + O(|\zeta|^3) \tag{52}$$

where $M^{(i)}$ is the ith derivative of $P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta}$ with respect to ζ, evaluated at $\zeta = 0$.

To obtain $M^{(1)}$, first differentiate using the product rule:

$$\tfrac{d}{d\zeta} P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta} = P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta}' + [P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta}']^{\scriptscriptstyle T}$$

Evaluating at $\zeta = 0$ gives,

$$M^{(1)} = P_0^{\scriptscriptstyle T} \Lambda \mathcal{E} + [P_0^{\scriptscriptstyle T} \Lambda \mathcal{E}]^{\scriptscriptstyle T}$$

Similarly, given the second derivative,

$$\tfrac{d^2}{d\zeta^2} P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta} = (P_{\zeta}'^{\scriptscriptstyle T} \Lambda P_{\zeta}' + P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta}'') + (P_{\zeta}'^{\scriptscriptstyle T} \Lambda P_{\zeta}' + P_{\zeta}^{\scriptscriptstyle T} \Lambda P_{\zeta}'')^{\scriptscriptstyle T}$$

we obtain

$$M^{(2)} = P_0^{\scriptscriptstyle T} \Lambda \mathcal{W} + 2\mathcal{E}^{\scriptscriptstyle T} \Lambda \mathcal{E} + \mathcal{W}^{\scriptscriptstyle T} \Lambda P_0$$

Substituting $\zeta_t = \zeta$ and $\Lambda_t = \Lambda$ in (52) gives the approximation,

$$\begin{aligned}\mathsf{E}[P_{\zeta_t}^T \Lambda_t^\Gamma P_{\zeta_t}] &= \mathsf{E}[P_0^T \Lambda_t^\Gamma P_0] \\ &\quad + \mathsf{E}[(P_0^T \Lambda_t^\Gamma \mathcal{E} + \mathcal{E}^T \Lambda_t^\Gamma P_0)\zeta_t] \\ &\quad + \tfrac{1}{2}\mathsf{E}[(P_0^T \Lambda_t^\Gamma \mathcal{W} + 2\mathcal{E}^T \Lambda_t^\Gamma \mathcal{E} + \mathcal{W}^T \Lambda_t^\Gamma P_0)\zeta_t^2] + O(\varepsilon^3)\end{aligned} \tag{53}$$

The first term on the RHS can be computed using Proposition 2.5. The second expectation involves $\mathsf{E}[\zeta_t \Gamma_t]$, which is computed in Lemma 2.6. For the third term we replace Γ_t by $\Gamma_t^\bullet + \widetilde{O}(\varepsilon)$, where $\Gamma_t^\bullet$ has mean π_0 and is independent of $\boldsymbol{\zeta}$. Combining all of these approximations gives the following approximation for the second term in (51):

$$\begin{aligned}\mathsf{E}[P_{\zeta_t}^T \Lambda_t^\Gamma P_{\zeta_t}] &= P_0^T \Pi_\varepsilon P_0 \\ &\quad + P_0^T \operatorname{diag}(\mathsf{E}[\Gamma_t \zeta_t])\mathcal{E} + \mathcal{E}^T \operatorname{diag}(\mathsf{E}[\Gamma_t \zeta_t])P_0 \\ &\quad + \tfrac{1}{2}(P_0^T \Pi_0 \mathcal{W} + 2\mathcal{E}^T \Pi_0 \mathcal{E} + \mathcal{W}^T \Pi_0 P_0)\mathsf{E}[\zeta_t^2] + O(\varepsilon^3)\end{aligned} \tag{54}$$

where, $\Pi_0 = \operatorname{diag}(\pi_0)$ and $\Pi_\varepsilon = \operatorname{diag}(\pi_\varepsilon)$. This gives (23) since $R_{\Gamma,\zeta}(0) = \mathsf{E}[\Gamma_t^T \zeta_t]$. □

E.2 Computation of $R_{B\zeta,\Delta}(t)$ in (26)

This is the most complex part of the proof.

We consider three cases separately: For $t < 0$ we have demonstrated that $R_{B\zeta,\Delta}(t) = 0$. The second case is $t = 0$. The approximation for $R_{B\zeta,\Delta}(0)$ is used as an initial condition in a recursive approximation for $R_{B\zeta,\Delta}(t)$, $t \geq 1$.

An approximation for $R_{B\zeta,\Delta}(0)$ is obtained from Lemma D.1. Using $\mathsf{E}[\Delta_0 | \mathcal{F}_{-1}] = 0$, we obtain

$$\begin{aligned}R_{B\zeta,\Delta}(0) &= \mathsf{E}[(\Gamma_0 \mathcal{E} \zeta_0)^T \Delta_0] \\ &= \mathcal{E}^T \mathsf{E}[\zeta_0 (\Gamma_{-1} P_{\zeta_{-1}} + \Delta_0)^T \Delta_0] \\ &= \mathcal{E}^T \mathsf{E}[\zeta_0 (\Gamma_{-1} P_{\zeta_{-1}})^T \Delta_0] + \mathcal{E}^T R_{\Delta^2,\zeta}(0) \\ &= \mathcal{E}^T \mathsf{E}[\zeta_0 (\Gamma_{-1} P_{\zeta_{-1}})^T \mathsf{E}[\Delta_0 | \mathcal{F}_{-1}]] + \mathcal{E}^T R_{\Delta^2,\zeta}(0) \\ &= \mathcal{E}^T R_{\Delta^2,\zeta}(0)\end{aligned} \tag{55}$$

Lemma D.1 provides an approximation for $R_{\Delta^2,\zeta}(0)$. Substituting this into (55) gives the approximation for $R_{B\zeta,\Delta}(0)$ shown in (26). □

In the remainder of this subsection we consider $t \geq 1$. Iterating (5) gives,

$$B_t^T = \Gamma_t \mathcal{E} = \Gamma_0 G_1 G_2 \cdots G_t \mathcal{E}$$

Each term is conditionally independent of Δ_0, given $\boldsymbol{\zeta}$, except for $\Gamma_0 = \Gamma_{-1} G_0 = \Gamma_{-1} P_{\zeta_{-1}} + \Delta_0$. Using the fact that $\Gamma_{-1} P_{\zeta_{-1}}$ is also conditionally independent of Δ_0, we obtain

$$\begin{aligned}R_{B\zeta,\Delta}(t) &= \mathsf{E}[\zeta_t \mathcal{E}^T P_{\zeta_{t-1}}^T P_{\zeta_{t-2}}^T \cdots P_{\zeta_0}^T [\Gamma_{-1} P_{\zeta_{-1}} + \Delta_0]^T \Delta_0] \\ &= \mathcal{E}^T \mathsf{E}[\zeta_t P_{\zeta_{t-1}}^T \cdots P_{\zeta_0}^T \mathsf{E}[\Delta_0^T \Delta_0 | \mathcal{F}_{-1}]] \\ &= \mathcal{E}^T \mathsf{E}[\zeta_t P_{\zeta_{t-1}}^T \cdots P_{\zeta_0}^T \mathcal{X}]\end{aligned} \tag{56}$$

where $\mathcal{X}$ was introduced in (44).

Denote $A_i = P_{\zeta_i}^T$, $A = P_0^T$, and the matrix product $\Omega_t = A_t A_{t-1} \cdots A_0$, for $t \geq 0$. We set $\Omega_t = I$ for $t < 0$. The product is approximated in the following:

Lemma E.1. *For $t \geq 0$,*

$$\Omega_t = A^{t+1} + \sum_{i=0}^{t} A^{t-i} \mathcal{E}^{\mathsf{T}} A^i \zeta_i + O(\varepsilon^2). \tag{57}$$

The proof of (57) is postponed to the end of this subsection.

Once we have established this lemma, we then have the complete cross-correlation:

Proof of approximation (26) *for* $R_{B\zeta,\Delta}(t)$, $t \geq 1$. We return to (56), which can be expressed

$$R_{B\zeta,\Delta}(t) = \mathcal{E}^{\mathsf{T}} \mathsf{E}[\zeta_t \Omega_{t-1} \mathcal{X}], \qquad t \geq 1.$$

Substituting the bound in Lemma E.1,

$$\begin{aligned}
R_{B\zeta,\Delta}(t) &= \mathcal{E}^{\mathsf{T}} \mathsf{E}[\zeta_t A^t \mathcal{X}] + \mathcal{E}^{\mathsf{T}} \mathsf{E}\Big[\zeta_t \sum_{i=0}^{t-1} A^{t-1-i} \mathcal{E}^{\mathsf{T}} A^i \zeta_i \mathcal{X}\Big] + O(\varepsilon^3) \\
&= \mathcal{E}^{\mathsf{T}} A^t \mathsf{E}[\zeta_t \mathcal{X}] + \mathcal{E}^{\mathsf{T}} \mathsf{E}\Big[\zeta_t \sum_{i=0}^{t-1} A^{t-1-i} \mathcal{E}^{\mathsf{T}} A^i \zeta_i \Big(\mathcal{X}^\bullet + \widetilde{O}(\varepsilon)\Big)\Big] + O(\varepsilon^3) \\
&= \mathcal{E}^{\mathsf{T}} A^t \mathsf{E}[\zeta_t \mathcal{X}] + \mathcal{E}^{\mathsf{T}} \sum_{i=0}^{t-1} A^{t-1-i} \mathcal{E}^{\mathsf{T}} A^i R_\zeta(t-i) \Sigma^{\Delta^\bullet} + O(\varepsilon^3)
\end{aligned}$$

This establishes (26) since $\mathsf{E}[\zeta_t \mathcal{X}] = R_{\Delta^2,\zeta}(-t)$. □

The proof of Lemma E.1 uses a Taylor series approximation for Ω_t:

$$\Omega_t = A_t \Omega_{t-1} = A\Omega_{t-1} + W_t\,,$$

where $W_t = E_t + F_t$; $E_t = \mathcal{E}^{\mathsf{T}} \zeta_t \Omega_{t-1}$ is the first order term in the Taylor series approximation, and F_t is the approximation error whose norm is bounded by $O(\varepsilon^2)$.

The following result provides a uniform bound on $A^{t-i} F_i$ for each t and i, and also $A^{t-i}\mathcal{E}^{\mathsf{T}}$ (which appears in (57)).

Lemma E.2. *The identity $\mathcal{E}1 = 0$ holds. Consequently, there exists $0 < \kappa < \infty$ and $0 < \varrho < 1$ such that*

$$\begin{aligned}
\|A^n \mathcal{E}^{\mathsf{T}}\| = \|\mathcal{E} P_0^n\| &\leq \kappa \varrho^n \\
\|A^n F_i\| &\leq \kappa \varrho^n \varepsilon^2
\end{aligned}$$

Proof. Since $P_\zeta 1 = 1$ for all ζ, we have

$$\mathcal{E}1 = \frac{d}{d\zeta} P_\zeta \Big|_{\zeta=0} 1 = \frac{d}{d\zeta} P_\zeta 1 \Big|_{\zeta=0} = \frac{d}{d\zeta} 1 = 0.$$

Next we apply the ergodic limit (19), recalling that $P_0^n \to 1 \otimes \pi_0$ geometrically fast as $n \to \infty$: there exists $0 < \kappa_0 < \infty$ and $0 < \varrho < 1$ such that

$$\|e_n\| \leq \kappa_0 \varrho^n, \quad \textit{with} \quad e_n = P_0^n - 1 \otimes \pi_0. \tag{58}$$

Consequently, $\mathcal{E} P_0^n = \mathcal{E} P_0^n 1 \otimes \pi_0 + \mathcal{E} e_n = \mathcal{E} e_n$, which implies the desired bound $\|\mathcal{E} e_n\| \leq \kappa \varrho^n$, with $\kappa = \kappa_0 \|\mathcal{E}\|$.

The proof of the second bound is similar: For each i we have

$$1^T W_i = 1^T (A_i - A)\Omega_{i-1} = 0^T$$
$$1^T E_i = 1^T \mathcal{E}^T \Omega_{i-1}\zeta_i = (\mathcal{E}1)^T \Omega_{i-1}\zeta_i = 0^T$$

We then have $1^T F_i = 1^T(W_i - E_i) = 0^T$, from which we obtain as before,

$$F_i^T P_\zeta^n = F_i^T (1 \otimes \pi_0 + e_n) = F_i^T e_n$$

Applying (58), we arrive at the desired bound:

$$\|A^n F_i\| = \|F_i^T P_0^n\| = \|F_i^T e^n\| \le \|F_i\|\,\|e^n\| \le \varepsilon^2 \kappa \varrho^n,$$

where $\varepsilon^2\kappa$ is equal to κ_0 times a worst-case bound on $\|F_i\|$. □

Proof of Lemma E.1. Applying Lemma E.2,

$$\sum_{i=1}^t \|A^{t-i}F_i\| = \sum_{i=1}^t \|F_i^T P_0^{t-i}\| \le \varepsilon^2 \frac{\kappa}{1-\varrho} = O(\varepsilon^2)$$

This implies the following approximation for Ω_t:

$$\begin{aligned}
\Omega_t = A_t\Omega_{t-1} &= (A + \mathcal{E}^T\zeta_t)\Omega_{t-1} + F_t \\
&= A\Omega_{t-1} + \mathcal{E}^T\Omega_{t-1}\zeta_t + F_t \\
&= A^t\Omega_0 + \sum_{i=1}^t A^{t-i}\mathcal{E}^T\Omega_{i-1}\zeta_i + \sum_{i=1}^t A^{t-i}F_i \\
&= A^t(A + \mathcal{E}^T\zeta_0 + O(\varepsilon^2)) \\
&\qquad + \sum_{i=1}^t A^{t-i}\mathcal{E}^T\Omega_{i-1}\zeta_i + O(\varepsilon^2) \\
&= A^{t+1} + \sum_{i=0}^t A^{t-i}\mathcal{E}^T\Omega_{i-1}\zeta_i + O(\varepsilon^2)
\end{aligned}$$

In particular, this shows that $\Omega_t = A^{t+1} + O(\varepsilon)$. Moreover, Lemma E.2 gives the geometric bound $\|A^{t-i}\mathcal{E}^T\| \le \kappa\varrho^{t-i}$. This justifies substitution $\Omega_{i-1}\zeta_i = A^i\zeta_i + O(\varepsilon^2)$ to obtain the desired result (57). □

E.3 Approximation of $R_{V\zeta^2,\Delta}(t)$ in (28)

Recall that $V_t^T = \frac{1}{2}\Gamma_t\mathcal{W}$, and denote $R_{\Gamma^\bullet,\Delta^\bullet}(t) = \mathsf{E}[(\Gamma_t^\bullet)^T\Delta_0^\bullet]$.

Applying the coupling result Proposition 2.3, the cross-correlation is approximated as follows:

$$\begin{aligned}
R_{V\zeta^2,\Delta}(t) &= \mathsf{E}[\tfrac{1}{2}(\Gamma_t\mathcal{W})^T\zeta_t^2\Delta_0] \\
&= \tfrac{1}{2}\mathcal{W}^T\mathsf{E}[\zeta_t^2]\mathsf{E}[(\Gamma_t^\bullet)^T\Delta_0^\bullet] + O(\varepsilon^3) \\
&= \tfrac{1}{2}\varepsilon^2\sigma_{\zeta^1}^2\mathcal{W}^T R_{\Gamma^\bullet,\Delta^\bullet}(t) + O(\varepsilon^3)
\end{aligned}$$

We have $R_{\Gamma^\bullet,\Delta^\bullet}(t) = 0$ for $t < 0$, and thus $R_{V\zeta^2,\Delta}(t) = O(\varepsilon^3)$ for $t < 0$.

We also have $R_{\Gamma^\bullet,\Delta^\bullet}(0) = \mathsf{E}[(\Gamma_0^\bullet)^T\Delta_0^\bullet] = \Sigma^{\Delta^\bullet}$, where $\Sigma^{\Delta^\bullet} = \Pi_0 - P_0^T\Pi_0 P_0$, with $\Pi_0 = \Pi_0$.

For $t \geq 1$,

$$\begin{aligned} R_{\Gamma^\bullet,\Delta^\bullet}(t) &= \mathsf{E}[(\Gamma_0^\bullet P_{\zeta_0} \dots P_{\zeta_{t-1}})^T \Delta_0^\bullet] \\ &= \mathsf{E}[((\Gamma_{-1}^\bullet P_{\zeta_{-1}} + \Delta_0^\bullet) P_{\zeta_0} \dots P_{\zeta_{t-1}})^T \Delta_0^\bullet] \\ &= \mathsf{E}[P_{\zeta_{t-1}}^T \dots P_{\zeta_0}^T (\Delta_0^\bullet)^T \Delta_0^\bullet] \\ &= \mathsf{E}[(P_0^T)^t (\Delta_0^\bullet)^T \Delta_0^\bullet] = (P_0^T)^t \Sigma^{\Delta^\bullet} . \end{aligned}$$

Substituting $R_{\Gamma^\bullet,\Delta^\bullet}(t) = (P_0^T)^t \Sigma^{\Delta^\bullet}$ into the previous expression for $R_{V\zeta^2,\Delta}(t)$ gives (28) for $t \geq 0$.

E.4 Proof of Proposition 2.8

The proof begins with the uniform bound,

$$\|\mathrm{S}(\theta) - \widehat{\mathrm{S}}(\theta)\| \leq \sum_{t=-\infty}^{\infty} \|\Sigma^{\mathrm{tot}}(t) - \widehat{\Sigma}^{\mathrm{tot}}(t)\|, \quad \theta \in \mathbb{R}$$

where $\|\cdot\|$ is any matrix norm. The right hand side is finite under Assumptions A1 and A2. It remains to obtain an estimate that is $O(\varepsilon^{2+\varrho})$. We establish a slightly stronger bound,

$$\sum_{t=0}^{\infty} \|\Sigma^{\mathrm{tot}}(t) - \widehat{\Sigma}^{\mathrm{tot}}(t)\| = O(\varepsilon^2 \log(1/\varepsilon)). \tag{59}$$

Under the assumption that $R_\zeta(t) \to 0$ geometrically fast, it follows that the same is true for $\widehat{\Sigma}^{\mathrm{tot}}(t)$ and $\Sigma^{\mathrm{tot}}(t)$: for some $b_0 < \infty$ and $\delta > 0$,

$$\|\Sigma^{\mathrm{tot}}(t)\| + \|\widehat{\Sigma}^{\mathrm{tot}}(t)\| \leq \exp(b_0 - \delta|t|), \qquad t \in \mathbb{Z}.$$

Moroever, $\mu = \mu_0 + O(\varepsilon^2)$, where the first d components of μ_0 coincide with π_0, and the remaining are zero. It follows that $\Sigma^{\mathrm{tot}}(t) = R(t) - \mu_0\mu_0^T + O(\varepsilon^4)$. Consequently, Theorem 2.4 implies that for some $b_1 < \infty$,

$$\|\Sigma^{\mathrm{tot}}(t) - \widehat{\Sigma}^{\mathrm{tot}}(t)\| \leq \exp(b_1)\varepsilon^3, \qquad t \in \mathbb{Z}.$$

To establish (59) we decompose the sum into two parts. Denote

$$N(\varepsilon) = \min\{t \geq 0 : \exp(b_0 - \delta t) \leq \exp(b_1)\varepsilon^3\}.$$

This implicit definition can be solved to give $N(\varepsilon) = [b_0 + 3b_1 \log(1/\varepsilon)]\delta^{-1}$.

From this we obtain,

$$\begin{aligned} \sum_{t=0}^{\infty} \|\Sigma^{\mathrm{tot}}(t) - \widehat{\Sigma}^{\mathrm{tot}}(t)\| &\leq \exp(b_1)\varepsilon^3 N(\varepsilon) + \sum_{t > N(\varepsilon)} \exp(b_0 - \delta t) \\ &\leq \exp(b_1)\varepsilon^3 N(\varepsilon) + \exp(b_0 - \delta N(\varepsilon)) \frac{1}{1 - \exp(-\delta)} \\ &\leq \exp(b_1)\varepsilon^3 N(\varepsilon) + \exp(b_1)\varepsilon^3 \frac{1}{1 - \exp(-\delta)} \end{aligned}$$

This together with the formula for $N(\varepsilon)$ immediately gives the bound (59). □

F Proof of Proposition 3.1

We first recall the representation of relative entropy as the convex dual of the log-moment generating function: For any probability measure ψ_b on $\mathcal{E}$ and measurable function $f\colon \mathcal{E} \to \mathbb{R}$, denote the log-moment generating function,

$$\Lambda_b(f) = \log(\psi_b(e^f))$$

For any other probability measure ψ_a on $(\mathcal{E}, \mathcal{B})$ we have,

$$\mathrm{D}(\psi_a \| \psi_b) = \sup\{\psi_a(f) - \Lambda_b(f)\}$$

where the supremum is over all measurable functions f for which $\psi_a(f)$ is defined [10, Theorem 3.1.2].

Provided the relative entropy is finite, the supremum is achieved uniquely with the log-likelihood function $f^* = \log(d\psi_a/d\psi_b)$. The error bound (ii) in the proposition is vacuous unless $\|f^*\|_\infty < \infty$. Consequently, to establish the error bound we can restrict to functions f for which $\|f\|_\infty < \infty$.

We apply the second order Taylor-series expansion,

$$\Lambda_b(f) = \log(1 + \psi_b(f) + \tfrac{1}{2}\psi_b(f^2)) = \psi_b(f) + \tfrac{1}{2}\psi_b(f^2) + O(\|f\|_\infty^3).$$

A quadratic approximation to relative entropy is obtained on dropping the error term. To complete the proof we establish the following alternate expression for (37):

$$\widehat{\mathrm{D}}(\psi_a \| \psi_b) := \sup\{\psi_a(f) - \psi_b(f) - \tfrac{1}{2}\psi_b(f^2)\}$$

where the supremum is over all functions f whose mean is defined with respect to both ψ_a and ψ_b. Without loss of generality we may assume that the maximum is over all functions f for which $\psi_b(f) = 0$. It is not difficult to show that the maximizing function is $\hat{f}^* = e^{f^*} - 1 = d\psi_a/d\psi_b - 1$ whenever $\widehat{\mathrm{D}}(\psi_a \| \psi_b)$ is finite. This establishes (i).

We can scale by a constant θ to obtain

$$\widehat{\mathrm{D}}(\psi_a \| \psi_b) = \max_f \max_\theta \{\theta\psi_a(f) - \tfrac{1}{2}\theta^2\psi_b(f^2)\}$$

where the first maximum is over measurable functions f satisfying $\psi_b(f) = 0$ and $\psi_a(|f|) + \psi_b(f^2) < \infty$. The optimizing value is $\theta^* = \psi_a(f)/\psi_b(f^2)$. Substitution leads to the formula (37). □